INVESTMENT PROCESS & ARTIFICIAL INTELLIGENCE

The Eight Verbs of Investing

Jai K. Jacob

Master chefs, we hear, conjure mouth-watering dishes from a handful of basics. Put those same ingredients in the wrong hands (er… mine), and you'll gain a newfound appreciation for reheated takeout. A bank-breaking trip to an artisanal grocer wouldn't elevate my output—but those foraged Matsutakes and Almas caviar turn into Michelin stars in skilled hands.

Our industry's fascination with markets, companies, trends, and forecasts—the raw ingredients of the investment meals we call portfolios—can distract us from the equally fascinating (if less emotionally riveting) topic of the investment craft itself. Just as innovations in agriculture and logistics changed how we eat, technology has transformed how we invest. From the spreadsheet to the database to the LLM, the speed at which we design and build portfolios has accelerated parabolically over the last thirty years. But what about *how* we invest?

A quant rarely opines on a specific investment or forecast (and thus rarely graces these hallowed pages) precisely because systematic investment concerns itself less with specific investments than with the *process* of investing. We've all heard about how to invest in AI. What about investing *with* AI?

AI—a catch-all term amalgamating machine learning (which includes LLMs), data science (which includes procurement and ingestion), and process automation (which includes agentification)—represents our industry's latest obsession. IT managers, freshly rebadged as AI experts, enjoy unconstrained budgets even as management fees compress and traditional research budgets shrink. The industry hopes, tacitly or explicitly, that AI will eventually bend the curve of labour unit costs downward over time, sparing Active Management's margins from a cost structure designed during an era of higher fees.

> **My claim, stated up front: every investment process, discretionary or systematic, reduces to eight verbs—Observe, Believe, Categorize, Qualify, Analyze, Rank, Weight, Commit. AI belongs to the verbs, not to the job titles. Any action that fails to accelerate or sharpen at least one of them probably doesn't warrant automating.**

Automate the process, not the job

In the early 1990s, we wrote a document on a computer, printed it, scanned it into a fax machine, digitally transmitted it to a remote printer, and scanned the result back onto a computer. That process automated the mail *courier*. Email automated the *delivery*. Technology should automate a process, not a job.

Investors don't manufacture product with PowerPoint funnels. Yet most AI initiatives in asset management automate couriers. To understand the process rather than the job, we must first abstract portfolio construction into concrete, evergreen actions. Verbs. AI shouldn't settle for automating existing jobs, constrained as those jobs remain by the physical world and human analytical capacity. The AI investment process thrives when we optimize it in digital space around the organizing principles of the verbs themselves.

01

Observe

We begin with *observe*. The composition and volume of observable data have changed profoundly. What began as structured tables (income statements, balance sheets) and physical interactions (company meetings, committee discussions) has exploded into video scraping, topical inference, satellite image parsing, social-media reaction analysis, and countless other measurable tidbits. To observe, in this sense, gathers; it does not process. If our ability to interpret complex data will grow over time, we should gather broadly today, even if some data remain unused at present. Today's digital exhaust might fuel tomorrow's intractable problem. Every interoffice chat, hallway conversation, half-read article, elevator wait time, and eye roll could yield tomorrow's interpretative context—and none of it need fit neatly into rows and columns.

The instinct to hoard indiscriminately, though, runs into real costs: storage, governance, regulatory exposure, downstream liability. The defensible stance treats observation as an active editorial choice, not a passive license to retain every keystroke.

Planet-scale infrastructure already exists. GDELT scans print, web, and broadcast media in 100+ languages on a 15-minute cadence, turning global news into structured signals. The 2024 Lowenstein Sandler alt-data survey found that 67% of hedge funds now employ alternative data, with AI cited as the key enabler for extracting signal from noise.

02 Believe

Next, we must *believe*. Observations interact—in our minds, in our algos—to form predictions, or hypotheses. A pattern we recognize or a trend we detect becomes a societal shift we anticipate or a line we extrapolate. To counter bias, structured beliefs identify confirming and undermining observations beforehand and specify kill conditions (among which should figure the passage of time). Beliefs, then, lean heavily on the quality and impartiality of our observations, and may soon expand past the domain of human reasoning entirely.

Consider a working hypothesis: sustained HNWI inflows, liberalized markets, and favourable tax zones will drive 2–3× growth in Gulf wealth management AuM by 2035, with UAE and Saudi Arabia as dual hubs. Qatar's HNWI population climbed roughly 560-fold over two decades; the UAE grew 20–40×. The hypothesis carries explicit KPIs (AuM, inbound UHNW flows, alt/credit penetration), decision gates (12/36/60-month AuM targets), and kill conditions (UAE HNWI inflows turn negative for two years; loss of free-zone tax benefits; geopolitical shocks cut flows by more than half). That structure separates a belief from a vibe.

BloombergGPT—a 50-billion-parameter model trained on roughly 700 billion tokens—outperforms general-purpose models on financial NLP tasks without losing general capability, lifting the ceiling for automated thesis generation across filings, transcripts, and Q&A. S&P Global Market Intelligence found that an LLM-based sentiment model applied to earnings calls would have delivered an 8.4% annualized long-short return—roughly double the 4.2% generated by traditional lexicon-based approaches.

03 Categorize

A broad belief often requires that we *categorize* it into tangible subsets. If we observe rising productivity in an economy and attribute it to automation, we next carve the problem: physical automation, digital automation, process improvement, and so on. An entire industry exists to classify broad markets into subsets—asset classes, sectors, industries, countries, themes—but when our belief inconveniently falls outside pre-existing buckets, we must allow specialists, human or machine, to decompose the broader problem and apply domain-specific expertise. Understanding the ontology of expertise already underlies successful agentic deployment; the adaptability of our approach materially improves resource allocation.

The GCC wealth thesis, for example, does not map cleanly onto GICS. It splits instead into seven driver-based categories: Prestige Wealth & Financial Services; Haute Lifestyle & Luxury Consumption; Elite Travel & Bespoke Mobility; Iconic Real Estate & Residences; Wellness, Aesthetics & Longevity; Elite Education & Cultural Experiences; Private Security & Lifestyle Management. Hoberg and Phillips's TNIC framework, built from 10-K product-description text and maintained through 2025 at Dartmouth Tuck, already demonstrates that text-based peer networks outperform SIC/NAICS in explaining competition, profitability, and manager-named peers.

04 Qualify

A portfolio does not consist simply of hypotheses; we must know which, among millions of tangible investments, our mandate permits. A *qualified* investment must align with our categorized beliefs and also clear scrutiny from regulators, benchmarks, risk managers, morals, operational capabilities, and client guidelines. Humans fall notoriously prey to cognitive traps here—the streetlight effect, the familiarity principle. Correcting those biases and building impartial eligible universes represents the most straightforward, immediate, and high-ROI application of data science to a fundamental investment process.

In the GCC case, the qualification funnel ran from roughly 1,278 candidate companies to 548 at least moderately aligned with the theme, narrowing further to 129 tightly aligned. The RUSBoost ensemble on raw accounting items (*Journal of Accounting Research*, March 2020) outperformed the Dechow logistic and SVM benchmarks by a wide margin in flagging likely restaters—practical for filtering higher-probability accounting risks before they reach an analyst's desk.

05 Analyze

Active management must discriminate among qualified investments. Each mode of research—factors to quants, styles to discretionary investors—applies domain-specific expertise to each qualified option, yielding some scalar measure of desirability: an alpha score, a price target, a buy/sell recommendation. Passive investment typically deploys a single mode (market capitalization); active managers deploy several—quality of management, valuation versus peers, underpriced growth, balance-sheet strength. Each offers a different angle on the same asset.

The 2019 MIT/ACM SIGMETRICS study showed that an ML model fusing anonymized weekly credit-card transactions with quarterly filings beat Wall Street analyst consensus on 57% of EPS predictions across 30+ firms. Consumer-spend exhaust, in other words, can carry alpha when the right ensemble fuses it with traditional inputs.

06

Rank

Modes disagree. A cheap-looking company may carry untrustworthy management; a fundamentally sound business may sit in a geopolitically unstable country. Reconciling these contradictions requires that, after we have analyzed each option across every pertinent mode, we *rank* all choices in the eligible universe from most to least desirable. Tradeoffs between modes—what quants call *factor loading*—can sit static ("valuation always drives 90% of the decision"), dynamic ("avoid politically unstable countries"), or domain-specific ("consult social media for consumer stocks"), and can scale in complexity up through subjective measures (a PM's familiarity with an analyst's biases) and full quantitative ensembles that rotate modes to fit recent conditions. At the end of ranking, we have sorted the universe by expected return.

Gu, Kelly, and Xiu's horse-race of ML models against classical regressions (*Review of Financial Studies*, 2020) showed that a value-weighted long-short decile strategy built on stock-level neural-network forecasts delivered an annualized out-of-sample Sharpe ratio of roughly 1.35—more than double the best regression-based benchmark. The gains came specifically from allowing nonlinear interactions that linear models miss.

07

Weight

Since portfolios contain multiple positions, the next phase considers interactions between highly-ranked choices and projects them across the plane of desired risk to *weight* investments. Rank preference alone can output a set too concentrated—or too diversified—to suit a portfolio. Two critical elements enter at this stage: a risk model—boundaries at aggregate and sub-aggregate levels (country, currency, position size, correlation), plus contingencies triggered by proximity or breach—and an objective function: maximize return, maximize reward-to-risk versus a benchmark, target a specified volatility level.

PMs may never explicitly define either, relying instead on intuition to size positions, weigh correlations, and trade off risk against reward. But whether intuitive or specified, the result emerges as a set of weights that incorporate—and ultimately define—the level of risk and return achievable by the portfolio as a whole.

Recent work on hybrid LSTM-GARCH architectures (Roszyk and Ślepaczuk, 2024) demonstrates meaningful improvement over vanilla GARCH for S&P 500 volatility forecasting—translating directly into sharper risk-scaling and tighter turnover control. In the GCC portfolio, the final weighting step allocates equal risk to each of the seven categories, a deliberately simple choice that still demands a correlation estimate to execute.

08 Commit

Once the portfolio holds investments and weights, the final stage *commits* those decisions to the market. Of all eight verbs, trading has seen the most technological advancement in public markets. The seniormost cohort of today's investment community recalls trading with paper tickets and animated phone conversations over dedicated landlines. Public-market execution has since shifted to algorithmic routing—VWAP, POV, implementation shortfall—monitored by real-time TCA and increasingly handled by execution copilots (Bloomberg EMSX, JP Morgan's LOXM) that switch algos mid-order and feed realized outcomes back into the research stack.

Private markets, though, remain largely untouched at the commit phase. The absence of observable, objective, actionable price discovery, along with the non-conforming nature of bespoke financing, has served both as a critical drawback and a prized feature of private markets. AI's ability to apply analytic frameworks to unstructured data holds promise for the asset class, though buyer and seller may privately agree that less transparency suits them both. Expect AI to compress diligence, comps, and structuring; do not expect it to erase opacity. Opacity, in this asset class, forms part of the product.

The GCC thesis, end to end

Stitched together, the eight verbs produce a single trace of an actual investment process:

Observe	Wealth distribution broadens sharply across the GCC.
Believe	WM services and ecosystem will grow 2–3× by 2035; UAE and Saudi Arabia as dual hubs, with defined KPIs and kill conditions.
Categorize	Seven driver-based categories, constructed from embeddings rather than GICS.
Qualify	1,278 candidates → 548 moderately aligned → 129 tightly aligned.
Analyze	93-name shortlist, scored on market cap and revenue purity to the theme.
Rank	By market-cap-weighted alignment score.
Weight	Equal risk across the seven categories.
Commit	Implemented via adaptive execution.

None of this demands AI. All of it benefits from AI—provided the benefit targets the verb, not the headcount.

What follows

Each verb of portfolio construction presents an opportunity for AI to knit together the techniques, practical learning, and domain expertise of discretionary investors with the breadth, scale, and risk precision of quantitative ones. Without the common language of these verbs, though, we risk deploying AI to automate tasks that would no longer need doing under a holistic process review that lifted the human constraint.

If every investment process reduces to these eight verbs, then any action inside one should measurably improve the efficiency or accuracy of at least one verb. Anything else automates a courier or polishes a PowerPoint.

Three implications follow. First, accountability does not migrate with the work. Source-grounded outputs, four-eyes review, audit logs, sandbox validation, and MNPI controls do not count as overhead—they earn AI outputs the right to enter a fiduciary process at all. Treat every AI claim that could shape a decision as non-actionable until it cites its source.

Second, any edge from AI inside a single verb will commoditize. The durable advantages live in proprietary data, tailored workflows, and the human judgment applied to question-framing and portfolio construction—the places where taste still matters.

Third, and most important: portfolio management does not die. It migrates—from craft to code, from intuition to inference, from discretion to design. The winners will architect accountable systems around the eight verbs. The losers will automate their couriers and call the result transformation.

REFERENCES: *Gu, Kelly & Xiu, "Empirical Asset Pricing via Machine Learning," Review of Financial Studies (2020); Hoberg & Phillips, Journal of Political Economy (2016) and updated TNIC data library (2025); Bao, Ke, Li, Yu & Zhang, "Detecting Accounting Fraud in Publicly Traded U.S. Firms Using a Machine Learning Approach," Journal of Accounting Research (March 2020); Roszyk & Ślepaczuk, arXiv (2024) on hybrid LSTM-GARCH; BloombergGPT (arXiv, March 2023); S&P Global Market Intelligence (2024); Lowenstein Sandler Alt-Data Survey (2024); GDELT Project; MIT/ACM SIGMETRICS (2019) on credit-card data and EPS prediction.*