Alphaxis: Snuffling for Truffles

Generated 2026-05-12 07:45 UTC

Alphaxis: Snuffling for Truffles

A Master Read on Pattern-Seeking in Markets to Generate Alpha

Compiled by Sherlock | *12 May 2026* | *Covering 27 books and papers*

OPENING: THE CORE QUESTION

Markets contain patterns. The question is whether those patterns are real, exploitable, and durable — or whether they are mirages conjured by selective memory, wishful mathematics, and the relentless urge of intelligent people to find order in noise.

That question has occupied the best quantitative minds in finance for over six decades. The literature it has produced runs from foundational papers that rewrote how we understand markets — Jegadeesh and Titman showing in 1993 that past winners keep winning, at least for a while — to devastating critiques suggesting that the vast majority of "discovered" patterns are nothing more than statistical artefacts: the inevitable output of testing 400 hypotheses on the same dataset and accepting whatever survives at the 5% level.

Both things can be true simultaneously. Some patterns are real. Most aren't. The job of a serious quantitative practitioner is to tell the difference — and then to act on the real ones before they decay.

Every book and paper in this reading list is trying, in one way or another, to answer a version of the same question: *where is the alpha hidden, and why hasn't it been arbitraged away?* The answers cluster into three broad camps.

The first camp is the **behaviouralists**. Their answer is that markets misprice assets systematically because human beings are irrational in predictable ways. We overreact to recent news. We underreact to earnings surprises. We hold losers too long and sell winners too early. We extrapolate trends that are already over. These biases create patterns that a disciplined, rules-based investor can exploit — not because the market is stupid, but because it is human. Kahneman, Shefrin, Statman, De Bondt, Thaler, and Shiller are the central voices here.

The second camp is the **factor hunters**. Their answer is that certain structural characteristics — size, value, momentum, profitability, low volatility — carry persistent return premia, whether because they represent genuine risk (the rational explanation) or because they are behavioural mispricings that limits-to-arbitrage prevent from being eliminated (the institutional explanation). Fama, French, Asness, Ilmanen, Berkin, and Swedroe lead this group.

The third camp is the **sceptics and methodologists**. Their answer, delivered with increasing force since 2015, is: be very careful. Harvey, Liu, and Zhu documented 316 published factors. Hou, Xue, and Zhang tried to replicate 452 anomalies and found 65% simply vanished under proper testing conditions. McLean and Pontiff showed that anomalies lose 58% of their returns post-publication as investors trade them away. The sceptics are not nihilists — they believe alpha exists — but they insist that the bar for believing in a pattern must be much, much higher than the academic literature has historically demanded.

Sitting above all three camps is a small number of practitioners who have actually done it at scale — Renaissance Technologies most famously, WorldQuant's Igor Tulchinsky at lesser-known but still extraordinary scale — whose existence proves that the truffle is real. Jim Simons and his mathematicians discovered that markets contain faint, recurring patterns across thousands of instruments, and that the edge in each pattern is tiny but real. The Medallion Fund compounded at 66% annually from 1988 to 2018. That is not luck. It is the result of extraordinary discipline in pattern discovery, validation, and execution.

The truffle is there. But most of what looks like a truffle isn't.

That is the intellectual landscape this reading covers. The Big Ideas below distil what the best minds in this field have learned — about where alpha hides, why it persists, how it dies, and what it takes to find the real thing before the crowd does.

THE BIG IDEAS

1. Momentum: The Most Robust Anomaly in Finance

If you had to bet on one pattern being real across markets, time periods, and asset classes, it would be momentum. The evidence is overwhelming and spans more than two centuries.

Jegadeesh and Titman's 1993 paper — still the most-cited work in this space — showed that buying the past year's top-performing stocks and shorting the worst performers generates approximately 1% per month in excess returns over 3-12 month holding periods. Crucially, this edge is not explained by market risk or common factors. It is something else.

What is it? Asness, Moskowitz, and Pedersen (2013) showed in their landmark "Value and Momentum Everywhere" paper that momentum works not just in US equities but across 40 countries, in bonds, currencies, and commodities. Moskowitz, Ooi, and Pedersen (2012) showed that time series momentum — where an asset's own past return predicts its future return — is equally persistent across 58 futures instruments. This universality is the strongest argument that momentum is a genuine feature of how information flows through markets, not a statistical artefact specific to one dataset.

The behavioural explanation is elegant: investors underreact initially (anchoring, conservatism bias), which causes the price to drift toward fair value rather than jump there immediately. The trend represents the market's slow, grinding update of expectations. The Gray and Vogel "Quantitative Momentum" framework distils this into practice: smooth momentum (steady rather than volatile upward progression) performs better than jagged momentum, and the quality of the trend path matters as much as its magnitude.

The risk: momentum is also the strategy with the most terrifying crash profile. It is net long beta in rising markets and net short in falling markets — exactly backwards. When volatility spikes and correlations rise, momentum strategies can lose 30-40% in a matter of weeks. This is the "momentum crash" that has culled many practitioners who failed to respect it.

The truffle hunter's lesson: Momentum is real, but it needs to be held with respect for its capacity to kill you. The edge requires fresh patterns — looking back 12 months but skipping the most recent month to avoid reversal noise — and it requires a regime filter that reduces exposure when the market is in sharp decline.

---

2. Value Works, But Slowly and Painfully

The value premium — cheap stocks outperforming expensive stocks — has roughly equivalent statistical power to momentum but operates over much longer timeframes and requires patience that most investors simply do not have.

Fama and French's three-factor model (1993) established size and value as systematic risk factors, explaining over 90% of portfolio return variation. Their later five-factor model added profitability and investment. The theoretical argument is that value stocks are genuinely riskier — they are distressed companies, vulnerable in bad economic times — and the premium compensates for that risk. The behavioural argument, pushed by De

Bondt and Thaler (1985), is that investors overreact to bad news, causing stocks to become cheaper than their fundamentals justify, and the subsequent mean reversion is the return.

Both arguments have merit. What's clear is that value investing requires extended holding periods (3-5 years), can underperform for a decade (as it did from 2009-2020), and is psychologically devastating to implement. Berkin and Swedroe's "Your Complete Guide to Factor-Based Investing" gives value high marks for persistence and pervasiveness, but notes correctly that the factor requires both a long time horizon and the institutional durability to survive the drawdowns.

Novy-Marx (2013) added an important nuance: gross profitability — revenue minus cost of goods sold, divided by assets — has as much predictive power as book-to-market value. In other words, quality is the other side of the value coin. Cheap and profitable is the sweet spot. This is why Fama and French eventually added profitability and investment to their model: the original value factor is incomplete without it.

The truffle hunter's lesson: Value alone is not enough. The combination of value and profitability/quality is more reliable than either alone. And value's negative correlation with momentum (-0.50 to -0.60 across markets, per Asness et al.) means holding both together dramatically smooths returns without sacrificing expected return.

---

3. The Factor Zoo Problem: Most "Discoveries" Are False

This is the most important and most underappreciated insight in the entire literature.

Harvey, Liu, and Zhu (2016) catalogued 316 factors published in top academic finance journals through 2014. They argued that the conventional significance threshold (t-statistic > 2.0) is wildly insufficient when hundreds of hypotheses are being tested on the same data. The correct threshold, accounting for the cumulative testing, should be closer to 3.0 — and even then, many published factors are likely false discoveries.

Hou, Xue, and Zhang (2020) made this concrete. They attempted to replicate 452 published anomalies and found that 65% fail to survive proper testing with microcaps removed and value-weighting applied. Imposing a multiple-testing correction raises the failure rate to 82%. The trading frictions literature — short-term patterns supposedly caused by market frictions — was almost entirely wiped out: 96% failed to replicate.

McLean and Pontiff (2016) added the publication decay dimension: anomalies lose 26% of their returns out-of-sample, and 58% post-publication, as investors learn about them and trade them away. The half-life of a published anomaly is not long.

This creates a brutal epistemological problem for quants. If 82% of published patterns are false discoveries, and another 16% decay rapidly post-publication, what is actually left? The answer appears to be: a small number of robust, pervasive, theoretically-grounded factors — momentum, value, profitability, low volatility — and then a vast, shifting sea of micro-edges that require constant refresh and exceptional execution.

The truffle hunter's lesson: The Alphaxis standard for claiming a real edge should be much higher than passing a single backtest. Every candidate pattern needs: (a) a theoretically plausible mechanism, (b) out-of-sample testing in different time periods, (c) pervasiveness across instruments or markets, and (d) survival under transaction cost assumption. Without all four, it's probably noise.

---

4. Behavioural Mispricings Are Real But Not Easily Harvested

Kahneman's "Thinking, Fast and Slow" provides the psychological substrate beneath every market anomaly. System 1 thinking — fast, emotional, pattern-matching — dominates investor behaviour under uncertainty. The implications for markets are profound and well-documented.

The disposition effect (Shefrin and Statman, 1985; Odean, 1998) shows that investors sell winners too early and hold losers too long — the exact opposite of what momentum theory prescribes. This creates a mechanical inefficiency: price keeps drifting upward after good news because too many investors are trimming positions, and price keeps drifting downward after bad news because holders are refusing to realize losses. Post-earnings announcement drift (Bernard and Thomas, 1989) is the direct consequence: earnings surprises continue to generate abnormal returns for 60-90 days because the market literally refuses to update quickly.

Shleifer and Vishny (1997) identified the institutional reason why these mispricings persist: the "limits of arbitrage." Real-world arbitrageurs use other people's money, face redemption risk, and cannot hold positions indefinitely against temporary adverse price movements. When a mispricing exists, the rational arbitrageur is constrained from correcting it by the fear of losing capital before the price reverts. This is why the same anomalies have persisted for decades: it's not that nobody sees them, it's that exploiting them at scale is genuinely difficult.

The truffle hunter's lesson: The most durable edges are anchored in persistent human psychology — and made available by structural constraints on arbitrage. PEAD survives partly because the arbitrage requires short positions, which are expensive and capacity-constrained. Understanding *why* a pattern exists is the best defence against it decaying.

---

5. Machine Learning Finds Real Patterns — But Has a False Discovery Problem of Its Own

The application of machine learning to financial prediction has shifted from naïve in the early 2010s to genuinely productive by the 2020s, but with important caveats.

Gu, Kelly, and Xiu (2020) tested the full machinery of modern ML — neural networks, gradient-boosted trees, penalized regressions — against a large cross-section of US stocks with 94 predictive characteristics. The results were striking: neural networks and tree-based methods significantly outperformed linear regressions, doubling out-of-sample Sharpe ratios in some specifications. The most important predictive signals were variations on momentum, liquidity, and volatility — the same factors that conventional quants had identified. ML wasn't finding entirely new patterns; it was finding more sophisticated *interactions* between known patterns.

Lopez de Prado's "Advances in Financial Machine Learning" (2018) provides the essential methodological corrective. His central argument is that most financial ML failures come from bad data handling and bad testing, not bad models. He identifies the key failure modes: using consecutive time-window cross-validation (which leaks future information), ignoring the non-stationarity of financial data, failing to account for the fact that financial observations are not independent (overlapping returns). His proposed remedies — purged k-fold cross-validation, combinatorial purged cross-validation, the triple barrier method for labelling — are now standard in serious research.

The honest verdict is that ML is a genuine tool for alpha discovery, but its false discovery problem is even worse than traditional statistics, because ML models have so many degrees of freedom that they can fit any pattern in the training data — including random noise. The truffle of genuine ML alpha is surrounded by a particularly deep undergrowth of spurious results.

The truffle hunter's lesson: ML is valuable for finding nonlinear interactions between known factors, not for discovering entirely new patterns. The validation methodology matters more than the model architecture. If you can't explain why a pattern exists, be very suspicious of a machine that found it.

---

6. Structural Patterns Are More Durable Than Statistical Ones

The most reliably persistent patterns in markets are not statistical — they are structural. They arise from the mechanics of how markets are organised, not from behavioural biases that might be arbitraged away.

Calendar effects are the classic example. The January effect — especially strong in small caps — persists because year-end tax-loss selling depresses prices artificially in December, which then reverse in January. The mechanism is structural (tax law), not behavioural. Similarly, options expiry dynamics create systematic price magnetism around large open interest strikes, because market makers' delta-hedging creates predictable flows. The CBOE's research on 0DTE options shows that dealers' gamma exposure can pin prices near key levels in the final trading hours — a structural feature of the options market, not a behavioural one.

PEAD is a hybrid: it persists partly because of behavioural underreaction, but also because exploiting it requires shorting — which is expensive and capacity-constrained. The structural cost of the trade preserves the anomaly.

Narang's "Inside the Black Box" identifies execution as itself a source of structural advantage: a quant firm that understands market microstructure and executes intelligently will capture more of its theoretical alpha than one that treats execution as an afterthought. The market impact of trades is a cost that destroys edge at scale — and understanding it is as important as finding the edge in the first place.

The truffle hunter's lesson: When assessing the durability of a pattern, always ask: what is the structural reason this anomaly persists? If the answer is "behavioural," it may decay as awareness grows. If the answer is "structural" — rooted in regulation, market mechanics, or institutional constraints — it is more likely to survive.

---

7. Pattern Decay Is Inevitable; The Question Is Speed

Nothing lasts forever. The most important practical implication of the academic literature is that alpha has a finite half-life, and managing that decay is as important as finding the alpha in the first place.

McLean and Pontiff (2016) quantified the publication decay: 58% of anomaly alpha disappears after a paper is published, because institutional investors read academic research and trade it. Recent mathematical modelling (2025) derives a hyperbolic decay function: $\alpha(t) = K/(1+\lambda t)$, where the rate of decay depends on how crowded the strategy becomes. Momentum's decay is slower than value's because momentum strategies reinforce each other — crowded momentum tends to sustain trends rather than collapse them. But eventually, even momentum strategies become over-crowded and crash.

The alpha decay framework from Maven Securities and others shows that individual signal alpha decays from roughly 54 basis points per month at signal formation to near zero by month 10-11 for typical equity signals. Diversification across many uncorrelated signals is the only sustainable defence.

Covel's "Trend Following" captures the practitioner's version of this insight: the only way to stay ahead of decay is to trade systems, not tips. Systematic traders who obey rules without override will, over time, avoid the cognitive biases that erode performance — but they also need to accept that any given rule will stop working eventually, and they need to keep innovating.

The truffle hunter's lesson: The Medallion Fund's secret is not one great pattern. It is a vast library of small patterns, each contributing a little, diversified across assets, constantly refreshed. Alphaxis needs a similar philosophy at its own scale: no strategy should ever be the only bet.

---

8. The Quant Infrastructure Is Part of the Edge

This is an insight that is conspicuously absent from the academic literature but very present in the practitioner canon.

Narang's "Inside the Black Box" describes the anatomy of a successful quant firm: alpha models, risk models, transaction cost models, portfolio construction, and execution — and insists that failure at any one of these layers destroys the edge generated by the others. A beautiful alpha model paired with naive execution is not a viable business.

Carver's "Systematic Trading" provides the practitioner framework for translating signals into positions in a way that is robust, diversified, and sized appropriately. His core contribution is showing that the decision about *how much to trade* is often more important than the decision about *what to trade*. Volatility targeting — sizing positions based on expected volatility to maintain constant risk exposure — is the foundational rule, and it is empirically superior to fixed position sizing.

Lopez de Prado adds the machine learning perspective: the biggest gains from ML in finance come not from prediction but from the surrounding infrastructure — better labelling of training data, better feature construction, better backtesting methodology, better transaction cost modelling. Improving the infrastructure around a mediocre signal can produce better live results than finding a great signal and running it through a naive framework.

The WorldQuant approach in "Finding Alphas" (Tulchinsky, 2015) takes this to its logical conclusion: an alpha is not a strategy — it is an individual signal. A portfolio of thousands of diverse, weakly-correlated alphas is more robust than any single strategy. Diversity of signals, not magnitude of any single signal, is the ultimate source of edge.

The truffle hunter's lesson: The truffle-hunting apparatus matters as much as the truffle itself. Poor execution, poor position sizing, and poor risk management will eat your alpha even if you find it.

---

9. The Very Best Evidence Points to Extreme Difficulty — and Extreme Reward

The Medallion Fund is the most important single data point in this entire literature. Zuckerman's "The Man Who Solved the Market" documents what happened when the world's best mathematicians and scientists applied themselves to pattern discovery in markets with unlimited data, institutional execution, and a culture of radical secrecy.

The result was 66% average annual returns before fees (approximately 39% after), sustained for 30+ years across all market conditions. No other investment vehicle in history comes close. The mathematical edge in each individual pattern was tiny — faint signals that would not survive casual statistical scrutiny. What Simons discovered is that the signal-to-noise ratio in financial markets is very low, but it is not zero. And that a vast number of tiny edges, properly diversified and efficiently executed, produces something extraordinary.

The three-step methodology attributed to Medallion — identify anomalous patterns; verify they are statistically significant, consistent over time, and non-random; then see if there is a reasonable explanation for the pattern's existence — is deceptively simple. The "reasonable explanation" requirement is the key filter: it guards against pure data mining. But Medallion also relaxed this requirement for patterns with sufficiently strong statistical evidence, which is intellectually uncomfortable but empirically profitable.

The important corollary is what Medallion did not do: it did not publish its findings. McLean and Pontiff proved that publication destroys alpha. Simons understood this intuitively before the paper was written. The academic literature is, in a sense, a gift to practitioners who find patterns before they are published and a tax on practitioners who follow published research.

The truffle hunter's lesson: The edge is real but small. It is found through patient, systematic exploration of enormous amounts of data, validated with extreme rigour, and exploited quietly before it becomes crowded. Alphaxis is, at its scale, doing what Medallion does at a vastly larger scale. The philosophy is identical; the execution infrastructure is the differentiator.

---

10. The Honest Sceptic's Position

The literature converges on a view that is uncomfortable for those who want simple answers. Markets are not perfectly efficient — there are too many documented, replicated, economically-significant anomalies for that to be true. But they are not so inefficient that alpha is easy to find — the replication crisis proves that.

The honest position, which O'Shaughnessy arrived at after 40 years of data, Asness has articulated many times, and Ilmanen's "Expected Returns" makes its central conclusion, is this: a small number of factors — momentum, value, profitability, carry, low volatility — have delivered persistent risk-adjusted returns over long histories, across multiple geographies, and survive most sceptical tests. These factors should be the foundation. Around the foundation, there are structural anomalies (calendar effects, PEAD, options expiry dynamics) that can be layered carefully. And at the frontier, there are micro-edges found by ML and signal-discovery — small, fast-decaying, requiring constant renewal.

The biggest mistake in this business is to mistake luck for skill, noise for signal. The second biggest mistake is to give up because patterns decay and replace them with nothing. The productive middle ground — evidence-based, humble about what we know, disciplined about what we trade — is exactly where the best long-term practitioners live.

THE READING LIST WITH VERDICTS

#	Title	Author	Year	Type	Sherlock's Verdict	Alphaxis Relevance
1	Returns to Buying Winners and Selling Losers	Jegadeesh & Titman	1993	Paper	**Gold** — the founding document of momentum	Direct: every momentum strategy we run traces to this
2	Expected Returns	Antti Ilmanen	2011	Book	**Gold** — e ncyclopaedic and rigorous	The best single reference for why factors work and when they don't
3	Advances in Financial Machine Learning	Lopez de Prado	2018	Book	**Gold** — essential methodology for the modern quant	Fix our backtesting infrastructure against its checklist
4	The Man Who Solved the Market	Greg Zuckerman	2019	Book	**Gold** — the proof that it can be done	The philosophy behind Medallion is the philosophy of Alphaxis at our scale
5	Value and Momentum Everywhere	Asness, Moskowitz, Pedersen	2013	Paper	**Gold** — proves universality of the core factors	Validates running momentum across crypto, not just equities
6	Empirical Asset Pricing via Machine Learning	Gu, Kelly & Xiu	2020	Paper	**Gold** — the definitive modern ML asset pricing paper	Confirms ML adds value but through known signals, not mystery signals
7	Time Series Momentum	Moskowitz, Ooi, Pedersen	2012	Paper	**Gold** — the academic foundation for trend following	Directly relevant to Flying High and Dual MA strategies

#	Title	Author	Year	Type	Sherlock's Verdict	Alphaxis Relevance
8	...and the Cross-Section of Expected Returns	Harvey, Liu & Zhu	2016	Paper	**Gold** — the most important sceptical paper in the literature	Raises our bar for claiming a real edge
9	Your Complete Guide to Factor-Based Investing	Berkin & Swedroe	2016	Book	**Silver** — excellent practical framework	Useful filter: only 5 factors meet all criteria
10	Inside the Black Box	Rishi Narang	2013	Book	**Silver** — best architecture guide to quant systems	The 5-layer model (alpha/risk/TC/portfolio/execution) should be our operating framework
11	Quantitative Momentum	Gray & Vogel	2016	Book	**Silver** — practical momentum construction guide	Momentum path quality (smooth vs jagged) is actionable
12	Thinking, Fast and Slow	Daniel Kahneman	2011	Book	**Silver** — the psychological substrate of every anomaly	Explains why our edges exist and why they don't decay overnight
13	Systematic Trading	Robert Carver	2015	Book	**Silver** — excellent position sizing and portfolio construction	Volatility targeting approach directly applicable
14	Replicating Anomalies	Hou, Xue & Zhang	2020	Paper	**Silver** — companion sceptical paper to Harvey et al	Confirms: most published patterns are dead on arrival
15	Does Academic Research Destroy Stock Return Predictability?	McLean & Pontiff	2016	Paper	**Silver** — quantifies publication decay	Quantifies why we should research privately and not share

#	Title	Author	Year	Type	Sherlock's Verdict	Alphaxis Relevance
16	Does the Stock Market Overreact?	De Bondt & Thaler	1985	Paper	**Silver** — foundational mean reversion paper	Mean reversion over 3-5 years is the counterbalance to momentum
17	Finding Alphas	Tulchinsky / WorldQuant	2015	Book	**Silver** — industrial-scale alpha discovery perspective	The portfolio-of-signals framework is our north star
18	Post-Earnings Announcement Drift	Bernard & Thomas	1989	Paper	**Silver** — foundational PEAD paper	Direct relevance to Project PEAD
19	The Limits of Arbitrage	Shleifer & Vishny	1997	Paper	**Silver** — explains why anomalies persist	Understanding why edges survive is as important as finding them
20	The Other Side of Value (Gross Profitability)	Novy-Marx	2013	Paper	**Background** — important but narrow	Relevant if we ever build equity long/short factor strategies
21	What Works on Wall Street	O'Shaughnessy	2011	Book	**Background** — 40+ years of data, slightly dated	Historical validation; useful for calibrating backtest expectations
22	Fact, Fiction and Momentum Investing	Asness et al. (AQR)	2014	Paper	**Background** — good for refuting myths	The 10 myths section is useful for building conviction
23	Evidence-Based Technical Analysis	David Aronson	2006	Book	**Background** — important methodological work	Core message (test rigorously, account for data mining) already in our DNA

#	Title	Author	Year	Type	Sherlock's Verdict	Alphaxis Relevance
24	Irrational Exuberance	Robert Shiller	2000	Book	**Background** — importantbut equity-valuation focused	CAPE framework useful for regime context, not direct signal generation
25	Trading and Exchanges	Larry Harris	2003	Book	**Background** — essential microstructure reference	Useful for execution design and understanding order flow patterns
26	Trend Following	Michael Covel	2017	Book	**Background** — inspirational, not technical	Good for conviction in systematic rules-based approach
27	Lucky Factors	Campbell Harvey	2015	Paper	**Background** — companion to Harvey, Liu & Zhu	Reinforces the multiple testing problem; less actionable than the main paper

---

WHAT TO READ FIRST

If Tony has time to read only three of these, in this order:

1. "Expected Returns" by Antti Ilmanen (2011)

This is the most complete single text in the field. It covers every major factor and strategy style, explains the evidence behind each, honestly assesses when each style works and when it fails, and builds a framework for portfolio construction across styles. If you only read one book about why certain patterns generate alpha, this is it. Allow 8-10 hours. Read chapters on value, carry, momentum, and the section on combining styles first — skip the fixed income deep dives on the first pass.

2. "Advances in Financial Machine Learning" by Lopez de Prado (2018)

The most important methodology book for anyone running systematic strategies today. Read the chapters on financial data structures, labelling, cross-validation, and backtesting. Even if you never use ML, the framework for thinking about data leakage and overfitting applies to every backtest we run. The chapter on the combinatorial purged cross-validation is dense — read it twice. Allow 6-8 hours on the relevant chapters.

3. Harvey, Liu & Zhu (2016) "...and the Cross-Section of Expected Returns"

Download the PDF from SSRN. It is 60 pages. Read it as a corrective discipline — a reminder that the bar for claiming a real pattern is much higher than it might feel when a backtest looks good. The key table of 316 factors with their t-statistics is genuinely humbling. Allow 3-4 hours.

This trio takes Tony from "what are the real patterns" (Ilmanen) through "how do I avoid fooling myself" (Lopez de Prado) to "just how much fooling myself is possible" (Harvey et al.). Everything else in the list builds on these three.

WHAT ALPHAXIS SHOULD DO DIFFERENTLY

Based on a thorough reading of this literature, here is an honest assessment of where Alphaxis's current approach is strong and where it could improve.

What we are doing well:

Alphaxis is running systematic, rules-based strategies — not discretionary trading. This is the single most important methodological choice in the field. Carver, Covel, Chan, and virtually every serious practitioner converges on this conclusion: systematic outperforms discretionary because it eliminates the behavioural biases that Kahneman identified. Our momentum strategies (Flying High, Dual MA) are directly aligned with the most robust anomaly in the academic literature.

The commitment to out-of-sample testing (2018-2025 in-sample, live from 2026) is sound. The walk-forward validation framework (WFE) is exactly what Pardo prescribes and what Lopez de Prado's combinatorial purged CV would produce in spirit if not in letter.

Where we should do more:

Factor diversification. Our strategy stack is heavily momentum-weighted (Flying High, Dual MA, Weekly Opening Range). This is the right foundation, but the literature on value/momentum correlation (-0.50 to -0.60) shows that holding a contrarian or value-tilted strategy alongside momentum improves Sharpe without reducing expected return. A quality-value screen applied to the equity universe — Novy-Marx's gross profitability measure — would be a natural complement.

Signal library, not just strategies. WorldQuant's "Finding Alphas" describes a model where hundreds of individual signals are combined into a portfolio. Alphaxis is building full strategies, each with complex logic. This is not wrong, but it means our diversification is coarse. A signal-level framework — where each market insight (EMA crossover, volume breakout, gap-and-go) is treated as an alpha signal and combined in a portfolio — would improve risk-adjusted returns and reduce concentration.

Publication lag exploitation. McLean and Pontiff proved that anomaly returns decay post-publication. Alphaxis has access to the academic literature and can trade patterns *before* they become crowded. The PEAD work (Project E1) is a good example of exploiting a known academic anomaly that persists due to structural constraints. More of this thinking — identifying structural anomalies that are well-documented but capacity-constrained in the institutional market — would be productive.

Regime awareness. Several papers (Moskowitz et al. 2012, Asness et al. 2013) document that factor performance varies significantly with market conditions. Ilmanen shows that momentum performs poorly in market crashes and well in trending markets. Our strategies should have explicit regime conditioning: reduce momentum exposure in bear markets, increase mean-reversion exposure. The Dual MA strategy's EMA50/EMA15 regime filter is a step in this direction, but it could be made more systematic and data-driven.

Execution and transaction costs. Lopez de Prado is categorical: most financial ML failures are execution and infrastructure failures, not model failures. The same applies to systematic strategies. We should have explicit transaction cost models that we run every signal through, and position sizing that accounts for our expected market impact. At small scale this is less critical, but as AUM grows it becomes the primary constraint.

The honest challenge:

The literature's core finding — that most patterns are false discoveries — is a challenge to Alphaxis's research culture. Strat Gym generates many tests. Not all tests that show a positive Sharpe are real. We need a standing policy that every strategy must pass Harvey et al.'s implicit t-statistic threshold of 3.0 (not 2.0), must survive at least two different out-of-sample periods, must have a stated mechanism, and must show fee-adjusted performance with realistic transaction costs before it is considered live-ready.

Currently the bar for "paper trading" is lower than it should be. A strategy that backtests well is not evidence of a real edge — it is a hypothesis. Paper trading tests market fit. But only live trading, with real capital and real execution, settles the question definitively.

THE BOTTOM LINE

There is a truffle buried somewhere in every market. It is real. It generates real returns. The best evidence — Medallion's 66% annual return, compounded across 30 years — proves it beyond reasonable doubt. Something is there.

But the undergrowth between the truffle hunter and the truffle is extraordinarily dense, and most of it is deceptive. The 452 anomalies that Hou, Xue, and Zhang tried to replicate looked, in their original papers, exactly like the real thing. They had t-statistics above 2.0, they had theoretical stories, they had peer review. And 82% of them were false discoveries. The academic finance industry has, for decades, been producing a kind of truffle-hunting map that mostly leads to empty holes.

What distinguishes the real from the fake? The real patterns share certain properties: they persist across very long histories (momentum goes back to Victorian Britain); they appear in multiple asset classes without requiring asset-class-specific stories; they survive with proper transaction cost modelling; they decay slowly even as awareness of them grows; and they have at least one plausible psychological or structural mechanism that explains why the inefficiency is not immediately arbitraged away.

The behaviouralists gave us the *why*: human beings are predictably irrational, and their systematic biases — loss aversion, the disposition effect, earnings underreaction — leave a trail that a disciplined, rules-based system can follow. The factor hunters gave us the *what*: momentum, value, profitability, low volatility, carry. The sceptics gave us the discipline: none of this is as easy as it looks, most published results are noise, and the bar for conviction must be very high.

For Alphaxis, the practical synthesis is this. We are a truffle-hunting firm. Our competitive advantage is not access to superior data (though that matters), nor to superior technology (though that matters too). It is methodology: the discipline to test rigorously, to demand multiple-period confirmation, to look for mechanisms rather than patterns, and to resist the temptation to declare a result real before it has earned that status through live performance.

The Medallion Fund's three-step methodology — identify, validate, explain — is simple enough to memorise but hard enough to execute that most practitioners never manage it. Alphaxis is trying to execute it at a fraction of Medallion's scale, with a fraction of Medallion's resources, but with the same intellectual framework. That is the right aspiration.

The truffles are there. The question is always: is this what I'm smelling actually a truffle, or is it just the damp earth making its familiar promises? Stay disciplined. Stay sceptical. And when the real thing surfaces — you'll know it, because it survives everything you throw at it.

Sherlock | *Alphaxis Partners* | *12 May 2026*

Reading list covers: Jegadeesh & Titman (1993), Fama & French (1993), De Bondt & Thaler (1985), Bernard & Thomas (1989), Shleifer & Vishny (1997), Moskowitz/Ooi/Pedersen (2012), Asness/Moskowitz/Pedersen (2013), Harvey/Liu/Zhu (2016), McLean & Pontiff (2016), Gu/Kelly/Xiu (2020), Hou/Xue/Zhang (2020), Novy-Marx (2013); Ilmanen (2011), Lopez de Prado (2018), Kahneman (2011), Narang (2013), Carver (2015), Berkin & Swedroe (2016), Gray & Vogel (2016), Tulchinsky (2015), Aronson (2006), O'Shaughnessy (2011), Harris (2003), Shiller (2000), Covel (2017), Zuckerman (2019), Chan (2013)

This document was produced by the Don/Ron quantitative research system. All backtests use historical data only. Past performance is not indicative of future results. Not financial advice.