From Signal Research to Capital Use

An Evidence Permission Model for Systematic Investment Implementation

A practitioner framework for converting research evidence into monitored portfolio process

Renato Guerrieri
Guerrieri Capital Ltd
June 2026

Abstract

In systematic investing, the weak point is often the handover from a research artefact into live portfolio use. A signal can appear credible in a backtest while remaining unsuitable for implementation once data lineage, point-in-time availability, the definition of the forward outcome, transaction costs, liquidity, capacity, portfolio constraints, attribution and monitoring are examined. Backtests are necessary evidence, but they are not approval to allocate capital. They may reflect data snooping, hidden look-ahead, survivorship effects, excessive selection, simplified execution assumptions or exposure patterns that are not visible in headline simulated performance.

The approach proposed here is an evidence permission model for systematic investment research. It has five parts: four separate permissions; a taxonomy of signal states; a control chain from research to implementation; a monitoring and attribution loop; and a clear boundary between research evidence, implementation permission and capital permission. The practical point is straightforward. A signal should not drift from thesis to portfolio because a backtest looks attractive. It should move only after decision gates, rejection records, portfolio translation, sizing review, cost and capacity analysis, attribution design and monitoring arrangements have been made explicit.

The discussion is about method and governance in systematic investment implementation. It does not disclose proprietary models, holdings, live recommendations or performance claims, and should not be read as investment advice or an offer to invest.

Keywords: Systematic investing, signal validation, point-in-time data, backtesting, portfolio construction, implementation risk, attribution, model governance.

JEL Classification: G11, G12, G14, G17, C58.

Executive Summary

Systematic research becomes economically relevant only when it can survive the conditions under which capital is actually deployed. A research signal may be well motivated, statistically interesting and persuasive in a backtest, yet still fail as an investment process. The failure point is usually not the first idea. It is the permission decision: deciding whether the evidence is strong enough, clean enough and implementable enough to move towards live use.

The organising distinction is between evidence and permission. Research evidence asks whether a thesis appears to contain useful information under controlled historical testing. Implementation permission asks whether that evidence can survive portfolio construction, constraints, cost, liquidity, sizing and operational control. Capital permission asks whether the signal is authorised for live use within defined limits. These decisions are related, but they are not interchangeable. Treating them as one decision creates process risk.

Backtests remain necessary. They are not sufficient. A backtest is conditional on its data, labels, universe, timestamps, rebalance assumptions, execution assumptions, cost model, portfolio construction method and research path. Published work on data snooping, multiple testing and backtest overfitting shows why repeated trials against the same history can produce results that look stronger than they are (Lo and MacKinlay, 1990; White, 2000; Harvey, Liu and Zhu, 2016; Bailey et al., 2017).

Operationally, the contribution has five components: four separate permissions, a taxonomy of signal states, a control chain from research to implementation, a monitoring and attribution loop, and a disciplined boundary between research evidence, implementation permission and capital permission. Together, they turn the question from whether a backtest is attractive to whether the evidence deserves a defined next state.

The control chain is practical rather than decorative: research idea → data availability → PIT validation → signal specification → historical test → robustness checks → rejection or freeze decision → portfolio construction → sizing and constraints → implementation review → attribution → monitoring → research revision. The sequence is not meant to be a universal template for every strategy. Its purpose is to stop ideas becoming live processes through narrative momentum, internal familiarity or an attractive simulation alone.

Failed research is kept in the record rather than cleaned away. Rejected tests, blocked signals, weak variants and killed signals are part of the evidence base. A research archive containing only surviving ideas is not an institutional evidence archive. It is a selective memory device, and it invites the same weak ideas to be rediscovered later.

Live use is not outside research; it is where live evidence begins. Attribution has to exist before deployment, not after losses. Monitoring compares realised behaviour with the approved evidence file, including performance, drawdown, costs, liquidity, turnover, exposures, data quality, operational incidents and signal family contribution. When the evidence, data or implementation changes materially, a new version is opened and reviewed again.

For Guerrieri Capital Ltd, the institutional posture is implementation focused. The objective is not to advertise signals. It is to demonstrate a serious method for deciding which research evidence deserves permission to move towards live portfolio use, and under what constraints.

Introduction: Why the Gap Between Research and Implementation Matters

Systematic investing is often described through ideas, datasets, models and backtests. Those are visible artefacts. They are not the full investment process.

The more important question is whether a research thesis can be converted into an operating workflow that remains repeatable under portfolio, trading, risk, monitoring and governance constraints. That conversion is the gap between research and implementation. It is where many systematic processes lose discipline.

The gap matters because research work and portfolio work have different tolerances. Research tolerates exploration. A researcher can test a variant, inspect a result, change a lag, remove a field, alter a universe, adjust a label and test again. That flexibility is useful in discovery. It becomes dangerous when it carries into capital use without records, versioning and decision gates.

Portfolio process requires a different standard. It needs source lineage, public availability assumptions, point-in-time validation, frozen signal specifications, clear rejection evidence, cost and capacity analysis, sizing controls, exposure constraints, attribution logic, monitoring thresholds and revision governance. It also needs decision ownership. Someone must know what is approved, what is blocked, what is live, what is being watched and what has been killed.

In this framework, **evidence permission** means the decision to let research move beyond controlled testing. The permission process evaluates a research artefact, constrains it and either denies it, delays it, freezes it, moves it to shadow monitoring or approves it for a specific implementation use. The point is not to claim that every systematic strategy needs the same template. Strategies differ by asset class, holding period, data type, execution style, mandate and risk budget. The narrower claim is that systematic research benefits from a visible approval process, so that research evidence does not become capital use by default.

The rule is simple: no silent promotion from research artefact to portfolio process.

A signal may deserve research permission. It may not deserve evidence permission. A signal with evidence permission may not deserve implementation permission. A signal that passes implementation review may still not deserve capital permission at the proposed size. Each stage answers a different question.

Separating the stages matters because live capital exposes errors that backtests often hide. Data may arrive later than assumed. A vendor field may be restated. A signal may overlap with an existing family. Turnover may rise during stress. Borrow costs may make a short leg unusable. An optimiser may express the signal through sector exposure rather than stock selection. A profitable live period may come from beta, not the intended signal. A weak live period may still be consistent with the evidence file. Without attribution and monitoring, the process cannot tell the difference.

Exhibit 1: Control Chain from Research to Implementation

Purpose: Exhibit 1 sets out the control chain from research thesis to monitored portfolio process. It is designed to show that implementation is not a straight line from idea to backtest to capital, but a staged sequence with decision gates and legitimate stopping points.

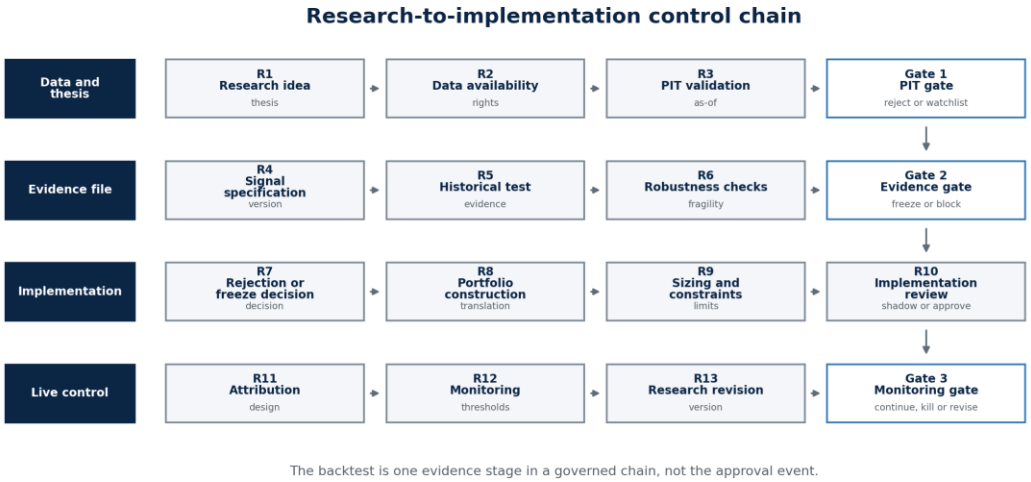

Caption. The control chain makes explicit where a research idea may be rejected, frozen, blocked, shadowed or revised before capital use. The backtest appears as one evidence stage within the chain, not as the approval event.

Contribution and Scope

The contribution is practitioner-facing rather than empirical. It does not introduce a new statistical test, pricing model, optimiser, execution algorithm or predictive signal. Its subject is the governance structure

around systematic investment implementation: how evidence is recorded, how decisions are separated and how capital use is controlled.

It has five components:

First, it separates four permissions that are often compressed into one informal judgement: research permission, evidence permission, implementation permission and capital permission. Each permission answers a different question and requires a different evidence standard.

Second, it defines a taxonomy of signal states covering idea, testing, rejected, frozen, watchlist, approved for implementation review, shadow, live, blocked and killed. The taxonomy treats rejection, blocking and killing as normal outcomes rather than failures to hide.

Third, it frames systematic implementation as a control chain rather than a research narrative. The chain begins with a thesis and ends with monitoring and research revision. It requires explicit decision gates before a signal enters portfolio use.

Fourth, it places portfolio construction between research evidence and capital permission. Portfolio construction is not a mechanical wrapper around a signal. It is the translation layer that determines whether the signal becomes the intended exposure or a different bet.

Fifth, it requires attribution and monitoring to be designed before deployment. The live process must know how it will judge whether the signal is working as intended, whether the portfolio is expressing unintended exposures, and whether the signal has drifted away from its approved specification.

The scope is institutional systematic research and implementation workflow. The argument applies most directly to systematic equity, quantamental, factor and cross-sectional processes, and to workflows driven by models or repeatable discretionary rules where research artefacts can drift into portfolio use without sufficient review and approval. Some concepts also apply to futures, macro, multi-asset and event-driven systematic processes, although the specific data, liquidity, execution and risk controls will differ.

No single control framework is sufficient for all strategies. The approach here does not prescribe a specific validation policy, regulatory compliance process, portfolio optimiser, execution method or risk budget. Nor does it address every asset-class-specific issue, such as derivatives margining, collateral management, market structure fragmentation, borrow recalls, futures roll methodology or OTC counterparty exposure. Those questions still require local judgement.

The intended use is narrower: provide a formal language and practical structure for deciding when research evidence is allowed to progress towards implementation, when it is denied, and how live behaviour is monitored after approval.

1. Research Artefacts Versus Investment Process

A research artefact is an output of investigation. It can be a signal definition, a notebook, a screening rule, a factor study, a model score, a test report, a chart pack, a research note or a code branch. It may be useful. It may be promising. It is still not an investment process.

An investment process is an operating system for decision-making. It specifies how inputs are sourced, when they are considered available, how signals are computed, how portfolios are formed, how trades are sized, how constraints are applied, how costs are estimated, how exceptions are handled, how performance is attributed and how live behaviour is monitored.

The distinction is frequently blurred in practice. A backtest is presented. A signal looks sensible. A portfolio construction variant appears acceptable. The signal is added to a model, initially at small weight. Over time, the fact that it is being used becomes the main justification for continued use. The original research state disappears.

That is not an implementation process. It is path dependence.

A disciplined process maintains clear answers to six questions:

What is the artefact? What decision has been made about it? What evidence supports that decision? What constraints apply? Who owns the next review? What would cause the state to change?

These questions are not administrative. They determine whether research is repeatable. A process that cannot reproduce its research output cannot govern that output. A process that cannot govern a signal is not ready to allocate capital to it.

Repeatability requires more than code storage. It requires the preservation of data versions, calendars, universe rules, lag rules, field definitions, corporate action handling, portfolio construction settings and cost assumptions. It also requires the preservation of failed work. Rejected variants are part of the research path. Removing them creates a cleaner narrative and a weaker evidence record.

Research artefacts also need context. A signal can be effective in raw rank form but not after sector control. It can work in large-cap names but not in the broader universe. It can look strong before transaction costs but fail after turnover. It can be statistically persuasive but economically duplicated by existing signal families. It can perform in a backtest because it carries beta, size or sector exposure rather than the intended information.

A portfolio process identifies these cases before live use. It also allows the answer to be “no”.

Institutional discipline is not measured by the number of signals approved. It is measured by the quality of evidence required for approval, the preservation of rejected evidence and the ability to identify when a live signal has moved away from its approved specification.

The working standard is therefore:

A research artefact may enter the evidence file. It may not enter the portfolio process without explicit permission.

2. The Evidence Permission Model

The four permissions separate decisions that are often collapsed into one approval. The structure is practical: it does not remove judgement, and it does not pretend that investment work can be reduced to a checklist. It makes the judgement visible, records the evidence supporting it and leaves a trace when the answer is no.

Permission	Core question	Typical evidence	Common failure if collapsed
Research permission	Is the idea worth testing?	Economic rationale, data availability, relevance to mandate, expected mechanism	Too many ideas are tested without thesis discipline, or too few are tested because the first data check is treated as final judgement
Evidence permission	Does the research evidence justify further review?	Point-in-time validation, frozen signal specification, historical test, robustness checks, rejected variants, sensitivity analysis	A backtest is treated as approval, even when evidence is fragile or selected after too many trials
Implementation permission	Can the signal be expressed inside the portfolio process?	Portfolio construction tests, exposure review, cost and capacity analysis, liquidity checks, sizing proposal, operational review	A signal with research merit becomes a different bet when constrained, sized or traded
Capital permission	Is the signal authorised for live use under defined limits?	Approved size, risk limits, monitoring thresholds, attribution setup, owner, review date, kill criteria	Live use begins without clear limits, monitoring logic or revision control

Research permission is deliberately light. It allows work to begin without implying belief in the signal. A research team needs room to test ideas before every early thesis is forced into an investment case. The control requirement is limited but important: the idea has a stated thesis, a data requirement and a disconfirmation condition before results start shaping the narrative.

Evidence permission is stricter. It asks whether the signal has survived enough controlled testing to deserve implementation review. This includes point-in-time validation, source lineage, forward outcome discipline, historical analysis, sensitivity to reasonable assumptions and a recorded account of rejected variants. Evidence permission is not capital permission.

Implementation permission tests the signal after translation into portfolio machinery. It asks whether the signal survives weighting, constraints, turnover control, liquidity limits, cost assumptions, risk model interaction and interaction with existing signals. A signal can pass evidence permission and fail implementation permission.

Capital permission is narrower still. It authorises use under specific conditions: size, universe, rebalance schedule, exposure limits, drawdown budget, monitoring thresholds, attribution logic, execution assumptions and review frequency. Capital permission is revocable. It is version-specific.

Exhibit 2: Evidence Permission Model

Purpose: Exhibit 2 shows the four permissions as distinct decision layers. The aim is to prevent a favourable backtest from being read as capital approval.

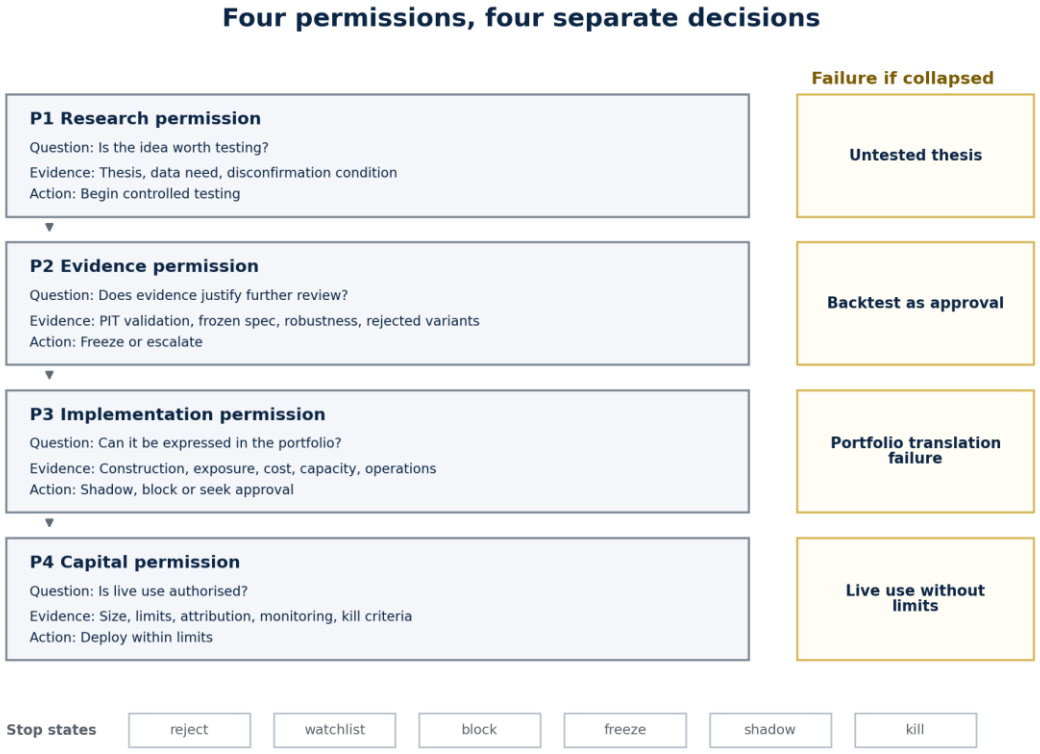

Caption. The four permissions separate the decision to test an idea, the decision to trust research evidence, the decision to assess implementation and the decision to allocate live capital under defined constraints. Each permission is narrower than the one before it.

Under this approach, the backtest stops being the central approval event. It becomes one item in a broader evidence file, alongside data controls, the original thesis, the signal definition, testing history, rejected variants, portfolio construction outcomes, cost and capacity assumptions, attribution setup and monitoring logic. A strong simulation can still matter, but it has to sit inside that wider record.

Failure also has a different meaning. A rejected signal is not wasted work if the reason is recorded. A blocked signal may be economically interesting but unusable because of data rights, operational limits or governance concerns. A frozen signal may be ready for shadow monitoring but not live use. A killed signal may have been valid under earlier conditions but no longer deserve capital permission.

Rejection, blocking, freezing and killing are therefore normal research outcomes, not embarrassments to be hidden. This matters because research processes weaken when only successful stories survive.

A formal permission record also preserves organisational memory. In many investment teams, the reason for a signal's existence is known by a small number of people. Over time, personnel change, code changes, vendors change and capital changes. Written decisions reduce dependence on informal memory and make later review less political.

None of this replaces investment judgement. The aim is to make judgement auditable, repeatable and less dependent on the most persuasive recent backtest.

3. Why Backtests Are Not Capital Permission

Backtests are necessary because systematic research needs controlled historical evidence. They are insufficient because the same historical evidence can be contaminated by data choices, repeated testing, implementation assumptions and hidden exposures.

The first problem is data snooping. If the same history is used repeatedly to select models, signals, transformations, labels, holding periods and universes, a strong result may reflect selection from many trials rather than a persistent relationship. Lo and MacKinlay (1990) and White (2000) formalise data-snooping concerns in asset pricing and model evaluation. Harvey, Liu and Zhu (2016) extend the concern to the large number of published return predictors, arguing that multiple testing changes the evidence hurdle for new factors.

The second problem is backtest overfitting. A research process can overfit not only through complex models, but through ordinary choices: the start date, lag assumption, universe filter, outlier rule, rebalance day, holding period, neutralisation method, cost level, signal blend and benchmark. Bailey et al. (2017) propose a framework for assessing the probability of backtest overfitting in investment simulations, a useful reminder that hold-out testing alone does not solve all selection risk when the broader research process is exploratory.

The third problem is implementation mismatch. A backtest may assume that trades occur at closing prices, that liquidity is sufficient, that borrow is available, that all positions can be sized, that constraints

do not bind and that turnover does not change during stress. These assumptions can be reasonable for a first test. They do not provide capital permission.

The fourth problem is exposure substitution. A backtest may appear to validate a signal while the portfolio return actually comes from beta, sector exposure, size, value, momentum, volatility, country allocation, currency, liquidity or concentration. If attribution is not designed before deployment, the process may not know whether the signal worked for the intended reason.

A backtest is conditional evidence. The condition set needs to be written down.

The evidence file records the thesis before the result is interpreted. It shows when data became available, how the forward label was defined, what universe was used, how corporate actions and delistings were handled, what variants were rejected, what costs were assumed, what exposures were present and what monitoring standard would apply if the signal were deployed.

A backtest can justify more work. It can justify rejection. It can justify freezing the signal for implementation review. It cannot, by itself, justify capital permission.

Vignette 1: When Timing Removes the Signal

A researcher tests a fundamental signal based on changes in reported balance-sheet strength. In the first test, the signal ranks companies using the latest available structured vendor field at each month-end. The backtest looks stable across several subperiods. Turnover is moderate, and the signal appears distinct from existing value and quality families.

The point-in-time review changes the conclusion. The field used in the test reflects a cleaned vendor record that was updated after the filing date. The simulated portfolio is therefore using information at the month-end rebalance that would not have been available in the production process. When the signal is rebuilt using conservative public-as-of timing, including filing delay and vendor ingestion delay, the apparent edge weakens materially. A further review shows that a large share of the original result came from the period between the simulated rebalance date and the realistic availability date.

The signal is not promoted. It is moved from testing to rejected, with the timing failure documented. The control lesson is not that fundamental data are unusable. It is that reported data must be tested against the decision timestamp, not against a historically accurate database value.

4. Data Discipline, Point-in-Time Control and Source Lineage

Data discipline begins before performance is examined. A signal using uncontrolled data is not an investment result. It is an unresolved data test.

Source lineage records where a field came from, when it was acquired, how it was transformed, how missing values were treated, how corrections were handled and whether vendor logic was applied. Lineage also identifies whether the field is raw, derived, manually adjusted, vendor-classified, internally modelled or produced by another upstream process.

Point-in-time control asks what the investment process could have known at the simulated decision time. This is not the same as asking whether the data value is historically true. A restated financial field may be true today. It may not have been available when the simulated portfolio traded. A classification may be correct today. It may not have been assigned at the time. An index member may appear in a current history. It may not have been part of the investable universe then.

Survivorship bias is one of the older examples of this problem. Brown, Goetzmann, Ibbotson and Ross (1992) show how truncation by survival can distort performance studies. Shumway (1997) documents delisting return issues in equity return data. Those concerns remain practical rather than historical: a backtest that excludes failed securities, stale listings or delisting events can overstate the investable opportunity set.

Public availability dates are a separate control. A filing date, public announcement time, vendor processing time, internal ingestion time and portfolio rebalance time may all differ. The conservative assumption is not always the most flattering assumption. It is the one that best reflects what the production process could have known.

Forward outcome discipline applies the same logic to outcomes. The predicted outcome is fixed before the signal is evaluated. If a signal is tested against several forward horizons, adjusted returns, benchmark-relative outcomes, volatility-scaled outcomes, event windows and exclusion rules until one form works, the process is no longer testing one signal. It is searching across labels.

Data entitlement belongs inside the control chain. A dataset may be available for research but not permitted for production use, redistribution, retention, derived-model use or audit archive. A signal blocked for data licence reasons may still be intellectually interesting. It is not implementable until the blocker is resolved.

Data controls also cover calendars. Reporting calendars, market holidays, half-days, exchange-specific closing times, timezone handling, corporate action effective dates and rebalance conventions can all change a simulated decision. Calendar errors are rarely dramatic. They can still determine whether a signal is real.

Point-in-time control is not a demand for perfect data. Perfect data do not exist. The standard is different: uncertainty is identified, treated conservatively and prevented from becoming hidden alpha.

A research process must be able to answer the following without reconstruction from memory: what data were used, when those data were assumed available, what fields were lagged, what fields were excluded, what gaps were filled, what corporate actions were applied, what universe was investable and what version of the dataset produced the result.

If those questions cannot be answered, the signal does not receive evidence permission.

5. Signal Specification, Rejection, Freeze and State Discipline

A signal needs to be specified tightly enough that another researcher can reproduce it without interpretation.

The specification includes the economic rationale, input fields, transformation rules, scaling method, ranking or scoring logic, lag assumptions, rebalance schedule, holding period, eligible universe, missing-data treatment, clipping or winsorisation, neutralisation, intended use and signal family. Intended use matters. A signal used as an alpha rank is not the same as a signal used as an exclusion filter, sizing input, risk control or timing overlay.

Signal family classification is central to implementation. Two signals may look different in code but express the same economic exposure. Several quality variants may select similar companies. Several sentiment signals may follow the same analyst revision cycle. Several defensive signals may load on beta, volatility and size. If the process treats such signals as independent, it may overstate diversification and understate crowding.

Signal decay deserves attention before approval. Decay can appear as weaker rank efficacy, shorter holding period, lower breadth, rising turnover, reduced spread, heavier exposure overlap or declining performance after publication. McLean and Pontiff (2016) document post-publication changes in return predictability for published predictors, while Hou, Xue and Zhang (2020) show that many published anomalies fail under stricter replication standards. These studies do not imply that all signals decay or fail. They support caution against assuming that historical evidence will survive research attention, capital pressure and implementation costs.

State discipline turns these concerns into an operating model. A signal does not exist only as “good”, “bad” or “used”. It has a state, an owner, a date, a version and a reason.

The required states are: idea, testing, rejected, frozen, watchlist, approved for implementation review, shadow, live, blocked and killed. The full taxonomy appears in Appendix A.

The most important state is often frozen. A frozen signal has a fixed definition for review. Once frozen, the signal is not improved casually because a later test looks better. Any material change creates a new version and reopens the evidence file.

The blocked state is also important. Some signals fail for reasons unrelated to historical performance. A signal may be blocked because data rights are insufficient, the data cannot be retained, the calculation cannot be reproduced, shorting is operationally unavailable, the signal depends on manual processing, the expected trades cannot be executed, or the proposed use creates disclosure or governance risk.

The shadow state allows a signal to be monitored without capital. This is useful when research evidence is promising but operational evidence is incomplete. Shadow monitoring can reveal data latency, file failures, unexpected turnover, unstable exposures, unrealistic execution timing and cost assumptions that do not survive production conditions.

The killed state is not a deletion. A killed signal leaves behind a record: why it was live, what changed, what monitoring evidence was observed, what alternatives were considered and who approved removal. Without this record, the same idea may return later under a new name with the same weakness.

Exhibit 3: Signal State Taxonomy

Purpose: Exhibit 3 shows the signal lifecycle as a governed state model rather than an informal set of research labels. It is intended to make rejected, blocked and killed states visible rather than hidden.

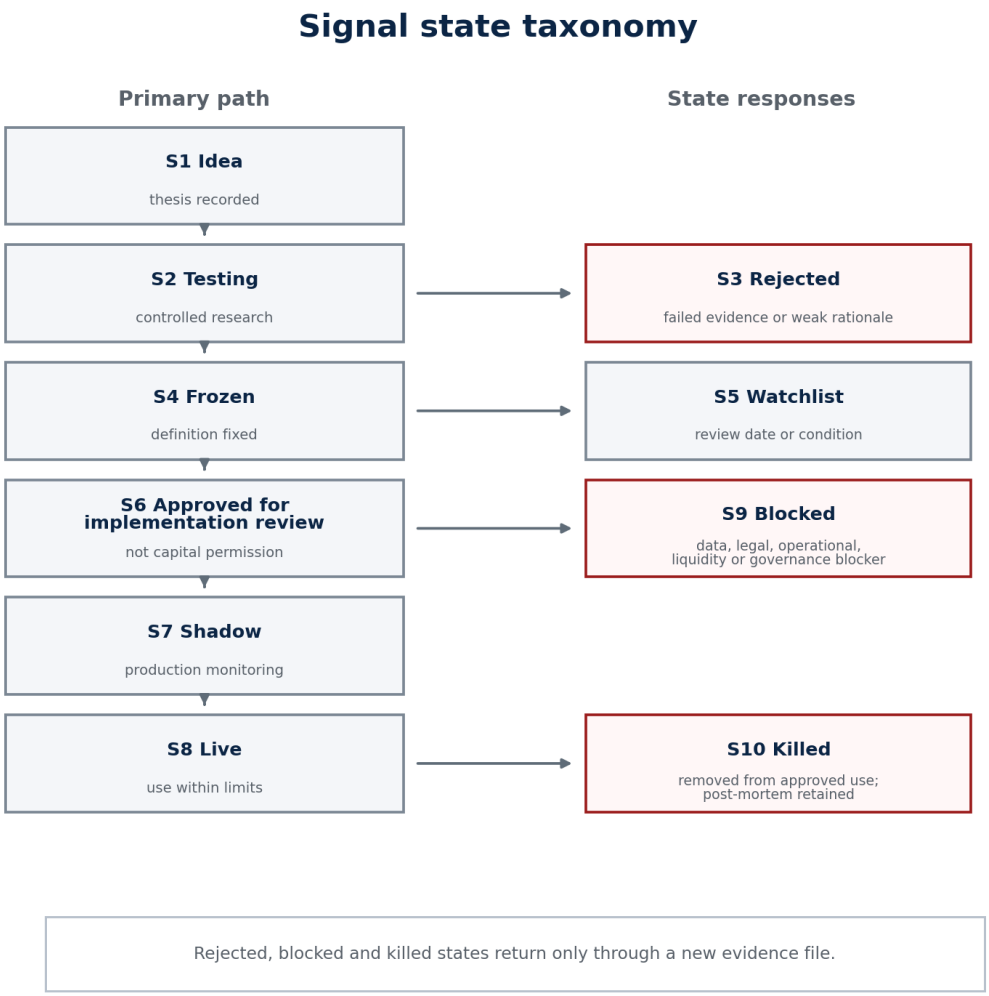

Caption. The state taxonomy prevents silent promotion and uncontrolled revision. Each state change requires evidence, ownership and a recorded decision, including decisions not to proceed.

Failed research is preserved, not hidden. Failed work reveals which data sources are weak, which families are crowded, which ideas are repeatedly selected by the same search process, which assumptions are fragile and which labels invite accidental look-ahead. It also reduces repeated research expense.

A research note is an evidence document. It is not written to sell the signal. It states what was tested, what failed, what changed, what remains uncertain and what decision was made.

The state model is the operational defence against silent promotion.

6. Portfolio Construction as the Translation Layer

Portfolio construction is where a signal becomes a capital expression. It is also where the signal can become something else.

A raw signal may rank names well in a historical test. Once converted into positions, it may require sector exposure, beta exposure, liquidity exposure, borrow availability, country concentration, factor tilts or a small number of names. The portfolio may not be expressing the original thesis. It may be expressing a construction artefact.

Implementation permission is separate from evidence permission for this reason. A signal with credible research evidence still has to pass through the intended portfolio machinery, where constraints, sizing, turnover and interactions with existing signals may change the conclusion.

The portfolio review covers weighting method, optimisation method, risk model interaction, turnover control, sector and factor constraints, name and issuer limits, liquidity rules, capacity assumptions, shorting rules where relevant, cash handling, rebalance frequency, execution timing and interaction with existing signals.

The same signal can behave differently under equal weighting, score weighting, risk weighting, volatility scaling, long-only construction, long-short construction or constrained optimisation. A signal may survive in an unconstrained spread test and disappear under sector neutrality. Another may look ordinary on its own but improve portfolio efficiency because it diversifies existing families. A third may have strong raw evidence but consume too much turnover or capacity.

Portfolio construction has long been treated as a central part of quantitative investment process, rather than a mechanical step after alpha research. Grinold and Kahn (2000) link expected return, risk and portfolio construction in their active management framework. Quantitative equity texts such as Qian, Hua and Sorensen (2007) treat construction, risk modelling and implementation as integrated components of equity process.

The translation layer is where the practical questions become unavoidable.

Is the portfolio taking the intended exposure? Are beta, sector, factor and concentration risks within limits? Is the signal contributing independently or duplicating an existing family? Does the optimiser behave sensibly, or is it producing unstable weights because constraints are tight? Is turnover consistent with the expected horizon? Can the names be traded at the proposed capital level? Does the drawdown budget match the mandate?

Drawdown budget is defined before capital permission. It is not a vague tolerance for loss. It specifies the measurement basis, review horizon, reference portfolio or benchmark, expected behaviour under stress and decision response if breached.

Sizing is part of the permission decision. A signal is not simply approved or rejected. It may be approved only for a small allocation, approved only within a restricted universe, approved only in shadow, approved only as a risk filter, or approved subject to additional monitoring.

Portfolio construction also tests whether the signal has value after constraint pressure. If a signal works only when allowed to take exposures outside the mandate, it is not suitable for that mandate. If it works only at trade sizes that cannot be executed, it is not suitable at the proposed capital level. If it works only when the optimiser is allowed to concentrate heavily, the evidence is concentration evidence, not broad signal evidence.

The implementation question is therefore not “does the signal work?” The better question is: **what portfolio does this signal create, and is that portfolio the one we intend to own?**

7. Costs, Capacity, Liquidity, Crowding and Regime Sensitivity

Implementation cost is part of the research result, not an adjustment made after a signal has already entered the portfolio narrative. If trading cost, liquidity or capacity consumes the edge, the historical signal has not survived implementation.

Transaction costs include more than commissions. A realistic implementation review may consider spread, taxes where applicable, financing, borrow, market impact, delay, participation rate, opportunity cost, slippage between decision price and execution price, and the cost of not completing trades. Perold (1988) remains a useful reference for framing the difference between paper decisions and executed outcomes through implementation shortfall. Almgren and Chriss (2001) formalise the trade-off between execution cost and execution risk in portfolio transactions.

Liquidity and turnover are assessed together. A universe can appear liquid while the strategy is hard to trade because positions rebalance at the same time, turnover clusters in smaller names, signal decay forces faster trading, or stress periods widen spreads and reduce depth. Liquidity is not only a property of a security. It is also a property of the trade list, capital base, participation rate, rebalance schedule and market condition.

Capacity is not a permanent number. It depends on universe depth, concentration, turnover, holding period, execution style, capital base, liquidity conditions, volatility and acceptable market impact. Novy-Marx and Velikov (2016) study anomaly performance after trading costs and cost-mitigation techniques.

Patton and Weller (2020) examine the gap between paper anomaly exposure and what implemented funds appear to capture. Frazzini, Israel and Moskowitz (2012) provide a practitioner perspective, based on institutional trading data, on trading costs and capacity. These strands do not imply one universal answer. They show that implementation cost can change the investment conclusion.

Crowding is harder to measure, but the absence of a clean measure does not justify ignoring it. Crowding can appear through factor overlap, common data sources, common portfolio constraints, similar rebalance schedules, correlated drawdowns, short-interest pressure, liquidity stress or exposure to well-known published predictors. Khandani and Lo (2011) analyse the August 2007 quant dislocation and provide a useful example of how similar systematic positioning and deleveraging pressure can create correlated stress. Arnott, Beck, Kalesnik and West (2016) discuss factor valuation and crowding concerns in smart beta strategies from a practitioner perspective.

Regime sensitivity is examined without pretending regimes can be known in advance. The useful implementation question is more modest and more practical: where did the signal historically fail, and what exposures explain those failures? Possible answers include liquidity stress, low dispersion, factor reversal, sector concentration, valuation compression, rising correlation, policy shock, financing pressure or data-cycle change.

Costs and regime sensitivity also connect to signal decay. A signal with slow decay and low turnover has more tolerance for imperfect execution than a fast-decay signal that must trade frequently. A short-horizon signal with modest gross edge may be highly sensitive to spread and impact. A signal that looks diversified in holdings may still be fragile if its edge depends on fast trading in less liquid names.

Vignette 2: The Signal That Could Not Pay for Trading

A cross-sectional signal is tested across a broad equity universe. It survives point-in-time review, appears stable across several subperiods and is not obviously explained by a single sector or factor exposure. On research evidence alone, it is a candidate for implementation review.

The portfolio test changes the conclusion. The signal decays quickly, so it requires frequent rebalancing. The highest-ranked names include many securities with moderate average volume and wider spreads. The long-only version creates turnover that is higher than the expected signal horizon can support. The market-neutral version is more expensive because both legs require active trading, and the short side introduces borrow and financing assumptions. When transaction costs are applied conservatively, the expected edge is consumed. Liquidity stress tests show that turnover would likely rise during the periods when the signal is weakest.

The signal is not killed because the thesis is false. It is rejected for capital use because implementation consumes the evidence. The research note preserves the result and records a possible future path: narrower universe, slower rebalance, or use as a secondary risk input rather than a primary alpha rank.

8. Attribution, Monitoring and Controlled Revision

Deployment is not the end of research. It is the start of live evidence collection.

Attribution is designed before a signal goes live. If attribution is built only after losses, it becomes forensic. It may still be useful, but it is late. A live process already knows how it will separate name contribution, sector contribution, factor contribution, signal family contribution, beta, concentration, turnover and cost effects.

Traditional performance attribution frameworks, including Brinson-style allocation and selection analysis, are not identical to systematic signal attribution, but they reinforce a basic principle: realised return is decomposed into explainable components rather than treated as a single outcome (Brinson, Hood and Beebower, 1986).

Systematic attribution is aligned with the evidence file. If the signal was approved as a quality signal, the monitoring process tests whether live returns are coming from the quality signal family or from unintended sector, beta, size, liquidity or country exposure. If the signal was approved as low-turnover, realised turnover is monitored. If capacity was approved under a participation assumption, realised execution is compared with that assumption.

Monitoring is not limited to performance. A profitable signal may be behaving incorrectly, while a losing signal may still be behaving as expected. The purpose of monitoring is to distinguish an adverse outcome from a process breach, because the two require different decisions.

Monitoring area	Review question	Possible state response
Performance	Is live return within the expected range of outcomes?	Continue, watchlist, review
Drawdown	Has the approved drawdown budget been breached?	Watchlist, reduce, freeze, kill
Turnover	Has realised turnover changed materially from the evidence file?	Review, revise, reduce
Cost	Are realised trading costs above assumptions?	Review, reduce, freeze
Liquidity	Are trades harder to execute than expected?	Reduce, block, revise
Exposure	Have beta, sector, factor or concentration exposures drifted?	Rebalance review, constrain, reduce
Attribution	Are returns coming from the intended signal family?	Continue, watchlist, revise
Decay	Is signal strength weakening relative to the frozen evidence file?	Watchlist, freeze, kill

Monitoring area	Review question	Possible state response
Data	Have source fields, coverage, vendor definitions or timestamps changed?	Block, freeze, retest
Operations	Have jobs, feeds, calendars, overrides or execution steps failed?	Incident review, block, revise
Governance	Is the live process still the approved version?	Freeze, reapprove, kill

Monitoring produces decisions, not only reports. A dashboard without state consequences is not a control system. The process defines what happens when thresholds are breached: continue, watchlist, freeze, reduce, block, kill or revise.

Material changes require version control. A live signal is not adjusted informally because recent results are uncomfortable. If the input fields change, the lag changes, the neutralisation changes, the universe changes, the weighting changes, or the signal family changes, the process creates a new version and reopens review.

Vignette 3: A Profitable Signal With the Wrong Source of Return

A signal is placed in shadow monitoring after passing research and implementation review. Over several months, the shadow portfolio appears profitable. A superficial review would support moving it to live status.

The attribution review gives a different answer. Most of the return is not coming from the intended signal family. The portfolio has developed an unintended sector tilt because the signal ranks a group of names favourably during a period when that sector outperforms. A secondary analysis shows that beta exposure also increased because the portfolio construction process allowed the signal to concentrate in higher-beta names within the sector. After sector and beta attribution, the residual signal contribution is modest.

The signal is not rejected immediately. It is returned to watchlist with a revised implementation question: does the signal retain value after sector and beta controls, and does the portfolio construction process express the intended information? The control lesson is that a profitable shadow period is not capital permission. Profitability has to be attributed before it is trusted.

Exhibit 4: Monitoring and Attribution Loop

Purpose: Exhibit 4 shows live or shadow monitoring as a feedback loop into state decisions and research revision. It prevents monitoring from being reduced to a reporting dashboard.

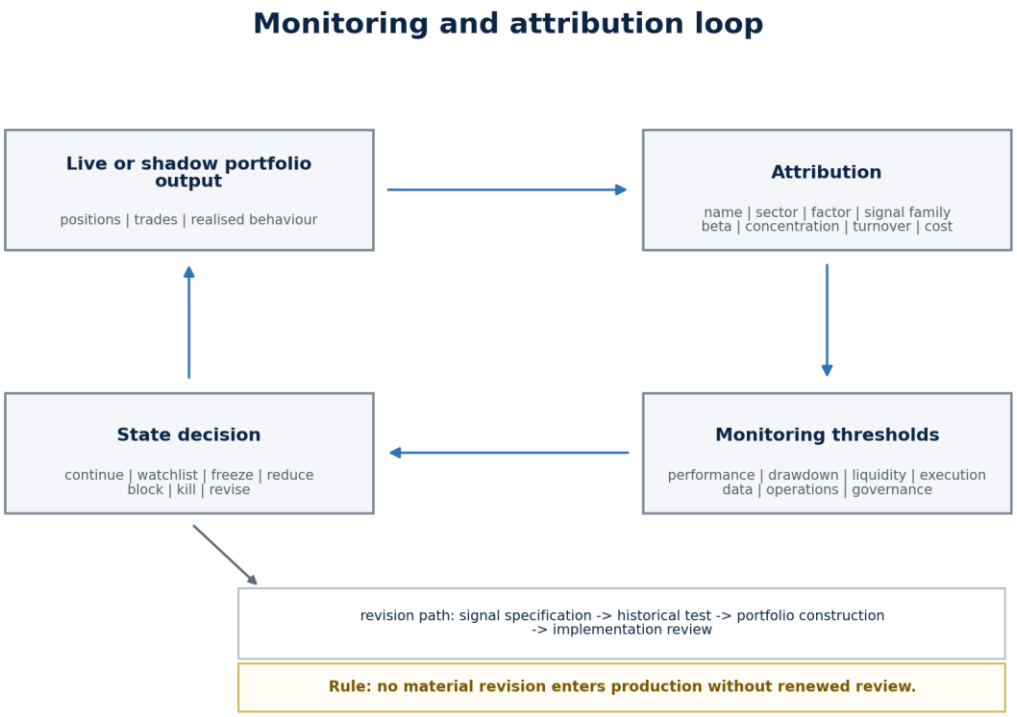

Caption. The loop links live evidence to attribution, monitoring thresholds and state decisions. It ensures that revisions pass back through controlled research and implementation review before entering the live process.

9. Public Evidence Boundaries and Disclosure Discipline

A private research vehicle distinguishes internal evidence from public communication.

Internal evidence may include signal definitions, rejected variants, proprietary data checks, code paths, data vendor details, cost models, capacity assumptions, portfolio constraints, monitoring thresholds and implementation decisions. These records are necessary for research discipline. They are not automatically suitable for public disclosure.

Public materials can explain process without implying product availability, capital raising, advisory services, performance history, client outcomes, holdings, live recommendations or forecast certainty. The discussion is methodological rather than product-specific. It is neither a description of an investable product nor a recommendation.

Disclosure discipline also applies to language. Research notes are evidence records, not promotional documents. They state what was tested, what failed, what changed, what remains uncertain and what

decision was made. If a note is later adapted for public use, confidential details are removed and factual claims are sourced.

Model governance standards in banking are not written for every systematic investment research process, but they show how financial institutions frame model risk through identification, governance, development, validation, use and monitoring. The Federal Reserve, OCC and FDIC issued revised model risk management guidance in 2026, replacing earlier guidance including SR 11-7 and the 2021 BSA/AML model-risk statement for the banking organisations in scope (Federal Reserve Board, OCC and FDIC, 2026). The PRA's SS1/23 sets out model risk management principles for relevant UK banks, building societies and PRA-designated investment firms using internal models for regulatory capital purposes (Prudential Regulation Authority, 2026). These sources are governance context, not a claim that the same regulatory requirements apply to Guerrieri Capital Ltd.

Public research also avoids selective presentation. If a paper discusses a framework, it does not imply that all internal signals pass the framework. If it discusses backtesting, it does not imply that simulated performance is live performance. If it discusses implementation, it does not imply that external capital is managed or sought.

The proper public posture is restrained. A process can be described without revealing proprietary signals. A research philosophy can be explained without presenting recommendations. An implementation framework can be set out without marketing performance.

The boundary is clear: internal evidence supports internal decisions; public writing explains methodology.

10. Practical Implementation Checklist

A practical checklist does not replace judgement. It makes judgement harder to bypass.

The checklist is used when a frozen signal is being considered for implementation review. It is not designed to make every early idea look institutional. Early research needs room for exploration; implementation review needs evidence, ownership and a decision record.

The minimum evidence file includes the original thesis, data lineage, point-in-time validation, forward outcome definition, signal specification, rejected variants, historical test, robustness checks, portfolio construction tests, cost and capacity analysis, sizing proposal, risk constraints, attribution plan, monitoring thresholds and revision rules.

A useful implementation review produces one of six decisions:

- Approve for live use under defined limits.

- Approve for shadow monitoring only.
- Approve conditionally, subject to specific blockers being resolved.
- Return to watchlist for further evidence.
- Block for data, legal, operational, liquidity, capacity or governance reasons.
- Reject.

The decision identifies the owner, version, date, approved use, limits, review schedule and criteria for escalation. It also specifies whether changes to data, code, universe, constraints, cost model or sizing require renewed review.

The research chain remains the practical backbone:

Research idea → data availability → PIT validation → signal specification → historical test → robustness checks → rejection or freeze decision → portfolio construction → sizing and constraints → implementation review → attribution → monitoring → research revision.

The chain is not paperwork after the fact. It shapes the order of work. Data permission precedes performance interpretation. Signal specification precedes final testing. Portfolio construction precedes capital permission. Attribution precedes deployment. Revision follows version control.

Appendix B provides a portfolio implementation checklist suitable for adaptation to an internal research memo, investment committee note or implementation review file.

The checklist carries one caution: a pass in every box does not guarantee success. It only states that the signal has been reviewed against the relevant controls. A systematic strategy can still fail because markets change, models are incomplete, data are flawed, costs rise, capacity falls, crowded trades unwind, or regimes shift.

The value of the checklist is not certainty. It is traceability.

Limitations

This is a procedural framework, and that is an important limitation. It does not guarantee investment performance, reduce uncertainty to an acceptable level in every case, or replace investment judgement. Its value is to state what evidence is required before a research artefact advances, what state the artefact occupies, and what monitoring applies after approval.

The approach deliberately does not depend on one method. It does not specify a particular alpha model, portfolio optimiser, risk model, transaction cost model, execution algorithm, statistical test or data vendor. That flexibility is necessary because systematic strategies differ by asset class, horizon, instrument, universe, liquidity, mandate and operating environment. The trade-off is that each

organisation must still define its own materiality thresholds, governance owners, escalation rules and documentation standards.

Nor does the framework solve the problem of unknown future market structure. A signal can pass point-in-time validation, cost review, capacity review and shadow monitoring, then fail because liquidity regimes change, a crowded exposure unwinds, data economics change, shorting conditions deteriorate, trading venues fragment, or policy shocks alter correlations. The permission process makes such risks visible and reviewable; it does not remove them.

Because the paper is intended for public release, it avoids proprietary formulas, holdings, securities, parameters, live ranks, track records, fund names and recommendations. That restriction limits the specificity of examples. The practitioner vignettes illustrate control lessons rather than actual strategy details.

The reference list is deliberately focused. It anchors the main implementation risks discussed here: data snooping, backtest overfitting, survivorship bias, trading costs, capacity, crowding, portfolio construction, attribution and model governance. This is a practitioner framework supported by selected sources, not a survey of the full academic literature.

Finally, the framework depends on honest records. A taxonomy and checklist can be misused if research teams preserve only favourable tests, relax thresholds after results are known, or treat sign-off as a formality. The structure is useful only if failed work remains searchable and state changes carry real consequences.

Conclusion

Systematic research earns institutional relevance only when it can pass from idea to implementation without losing evidence discipline.

The core failure is not the absence of backtests. It is the overuse of backtests as permission. A backtest is evidence under assumptions; it is not approval. By itself it does not answer whether the data were available, whether the signal is distinct, whether costs are tolerable, whether capacity exists, whether the portfolio expresses the intended bet, whether drawdown is acceptable, whether attribution is ready, or whether live revision will be controlled.

The evidence permission model addresses this failure by separating four decisions: research permission, evidence permission, implementation permission and capital permission. It adds state discipline through idea, testing, rejected, frozen, watchlist, approved for implementation review, shadow, live, blocked and killed. It keeps a control chain from research thesis to monitoring, treats attribution as infrastructure

that must exist before deployment, and requires material changes to create new versions and renewed review.

The point is not that every systematic strategy needs the same operational template. The point is that serious systematic research needs explicit review and approval. Ideas are tested, rejected, frozen, blocked, shadowed, approved, monitored and revised through a process that can be inspected.

The most important cultural implication is the preservation of failed research. Rejections, blockers and killed signals are not embarrassing artefacts. They are part of the evidence base. They prevent repeated errors, reveal weak data, expose families selected by repeated testing and strengthen future judgement.

Systematic implementation is not only a matter of models. It is a matter of permission, constraint and review.

Research evidence is not implementation permission.

Implementation permission is not capital permission.

No signal should cross those boundaries silently.

Author Note

Renato Guerrieri is the principal of Guerrieri Capital Ltd, a private research vehicle focused on systematic investment research and implementation infrastructure. The paper reflects an implementation focused view of systematic research, with emphasis on data control, portfolio construction, monitoring, attribution and evidence discipline.

Disclaimer

The material is provided for information and research purposes only. It does not constitute investment advice, investment research, a recommendation, an offer, a solicitation, or an invitation to buy or sell any security, financial instrument, investment product or investment service.

Nothing in this paper should be interpreted as a representation of current or future performance, a description of any investable product, or a commitment by Guerrieri Capital Ltd to manage capital, provide advisory services or make any investment available.

Examples are generic and suitable for public release. No holdings, tickers, securities, strategies, track records, forecasts, rankings, client outcomes or proprietary formulas are presented.

Systematic investment processes involve material risks, including model risk, data risk, implementation risk, liquidity risk, transaction cost risk, capacity risk, crowding risk and market risk. Simulated historical results, where used in research, are not a reliable indicator of future results.

AI assisted drafting and editing tools were used in preparing this manuscript. The author remains responsible for the paper's arguments, structure, review, factual accuracy and final content.

References

Almgren, R. and Chriss, N. (2001). "Optimal Execution of Portfolio Transactions." *The Journal of Risk*, 3(2), 5-39. DOI: 10.21314/JOR.2001.041.

Arnott, R. D., Beck, N., Kalesnik, V. and West, J. (2016). "How Can 'Smart Beta' Go Horribly Wrong?" *Research Affiliates*, February. DOI: 10.2139/ssrn.3040949.

Bailey, D. H., Borwein, J. M., López de Prado, M. and Zhu, Q. J. (2017). "The Probability of Backtest Overfitting." *The Journal of Computational Finance*, 20(4), 39-69. DOI: 10.21314/JCF.2016.322.

Brinson, G. P., Hood, L. R. and Beebower, G. L. (1986). "Determinants of Portfolio Performance." *Financial Analysts Journal*, 42(4), 39-44. DOI: 10.2469/faj.v42.n4.39.

Brown, S. J., Goetzmann, W. N., Ibbotson, R. G. and Ross, S. A. (1992). "Survivorship Bias in Performance Studies." *The Review of Financial Studies*, 5(4), 553-580. DOI: 10.1093/rfs/5.4.553.

Federal Reserve Board, Office of the Comptroller of the Currency and Federal Deposit Insurance Corporation. (2026). *Revised Guidance on Model Risk Management*. SR 26-2, OCC Bulletin 2026-13 and FDIC FIL-15-2026, 17 April 2026.

Frazzini, A., Israel, R. and Moskowitz, T. J. (2012). "Trading Costs of Asset Pricing Anomalies." Fama-Miller Working Paper, Chicago Booth Research Paper No. 14-05. DOI: 10.2139/ssrn.2294498.

Grinold, R. C. and Kahn, R. N. (2000). *Active Portfolio Management: A Quantitative Approach for Providing Superior Returns and Controlling Risk*. 2nd ed. New York: McGraw-Hill.

Harvey, C. R., Liu, Y. and Zhu, H. (2016). "... and the Cross-Section of Expected Returns." *The Review of Financial Studies*, 29(1), 5-68. DOI: 10.1093/rfs/hhv059.

Hou, K., Xue, C. and Zhang, L. (2020). "Replicating Anomalies." *The Review of Financial Studies*, 33(5), 2019-2133. DOI: 10.1093/rfs/hhy131.

Khandani, A. E. and Lo, A. W. (2011). “What Happened to the Quants in August 2007? Evidence from Factors and Transactions Data.” *Journal of Financial Markets*, 14(1), 1-46. DOI: 10.1016/j.finmar.2010.07.005.

Lo, A. W. and MacKinlay, A. C. (1990). “Data-Snooping Biases in Tests of Financial Asset Pricing Models.” *The Review of Financial Studies*, 3(3), 431-467. DOI: 10.1093/rfs/3.3.431.

McLean, R. D. and Pontiff, J. (2016). “Does Academic Research Destroy Stock Return Predictability?” *The Journal of Finance*, 71(1), 5-32. DOI: 10.1111/jofi.12365.

Novy-Marx, R. and Velikov, M. (2016). “A Taxonomy of Anomalies and Their Trading Costs.” *The Review of Financial Studies*, 29(1), 104-147. DOI: 10.1093/rfs/hhv063.

Patton, A. J. and Weller, B. M. (2020). “What You See Is Not What You Get: The Costs of Trading Market Anomalies.” *Journal of Financial Economics*, 137(2), 515-549. DOI: 10.1016/j.jfineco.2020.02.012.

Perold, A. F. (1988). “The Implementation Shortfall: Paper vs. Reality.” *The Journal of Portfolio Management*, 14(3), 4-9. DOI: 10.3905/jpm.1988.409150.

Prudential Regulation Authority. (2026). *Model risk management principles for banks*. Supervisory Statement SS1/23, current version published and effective 23 April 2026, originally published 17 May 2023. London: Bank of England.

Qian, E. E., Hua, R. H. and Sorensen, E. H. (2007). *Quantitative Equity Portfolio Management: Modern Techniques and Applications*. Boca Raton: Chapman & Hall/CRC.

Shumway, T. (1997). “The Delisting Bias in CRSP Data.” *The Journal of Finance*, 52(1), 327-340. DOI: 10.1111/j.1540-6261.1997.tb03818.x.

White, H. (2000). “A Reality Check for Data Snooping.” *Econometrica*, 68(5), 1097-1126. DOI: 10.1111/1468-0262.00152.

Appendix A: Signal State Taxonomy

State	Definition	Entry evidence	Permitted next states	Control note
Idea	A thesis has been identified but not yet tested.	Rationale, expected mechanism, required data, disconfirmation condition.	Testing, rejected, blocked.	The idea is recorded before results are known.
Testing	The signal is under controlled research review.	Data access, initial specification, outcome definition, initial tests.	Rejected, frozen, watchlist, blocked.	Testing variants remain preserved.
Rejected	Evidence does not justify further work or use.	Failed tests, weak rationale, data defect, exposure duplication or cost failure.	Idea or testing only if a new thesis or new evidence file is opened.	Rejection reasons remain searchable.
Frozen	Signal definition is fixed for review.	Final specification, data lineage, PIT validation, evidence file.	Approved for implementation review, watchlist, rejected, blocked.	Material changes require a new version.
Watchlist	Evidence is not sufficient for approval, but the idea remains relevant.	Unresolved concerns, need for more evidence from live conditions, sensitivity to assumptions.	Testing, frozen, rejected, blocked.	Watchlist status includes a review date or review condition.
Approved for implementation review	Research evidence is sufficient to test portfolio expression.	Frozen signal, historical evidence, rejected variants, data approval.	Shadow, live, watchlist, blocked, rejected.	This is not capital permission.
Shadow	Signal is monitored in production conditions without capital.	Implementation review, shadow portfolio rules, monitoring plan.	Live, watchlist, blocked, killed.	Shadow results require attribution before promotion.
Live	Signal is approved for portfolio use under defined limits.	Capital permission, sizing, constraints, attribution, monitoring thresholds.	Continue, watchlist, freeze, reduce, blocked, killed, revised version.	Live status is version-specific.
Blocked	Signal cannot be used due to data, legal, operational, liquidity, capacity or governance issue.	Blocker record, owner, remediation condition.	Testing, frozen or implementation review only after blocker resolution.	Blocked signals do not enter production through workarounds.
Killed	Signal is removed from approved use.	Monitoring evidence, review decision, post-mortem.	Idea or testing only through a new evidence file.	Killed signals remain in the archive.

Appendix B: Portfolio Implementation Checklist

Review area	Required question	Status options	Evidence to retain
Research thesis	Is the economic or behavioural rationale stated before final testing?	Pass, fail, conditional	Original thesis note and disconfirmation condition
Data lineage	Can each input be traced to source, ingestion and transformation history?	Pass, fail, conditional, blocked	Data map, vendor notes, transformation records
Data entitlement	Is production use permitted under the relevant licence or internal policy?	Pass, fail, conditional, blocked	Licence review or entitlement record
Public-as-of timing	Were inputs available before the simulated decision timestamp?	Pass, fail, conditional	PIT test, lag assumptions, availability calendar
Forward outcome	Was the predicted outcome fixed before signal selection?	Pass, fail, conditional	Outcome definition and version record
Universe	Is the investable universe free from survivorship and membership leakage?	Pass, fail, conditional	Universe construction file and historical membership logic
Signal freeze	Is the signal definition fixed for implementation review?	Pass, fail	Frozen specification, version number, owner
Rejected variants	Are failed variants and search-path decisions preserved?	Pass, fail, conditional	Research archive and rejection notes
Historical evidence	Does the frozen signal survive the required historical tests?	Pass, fail, conditional	Test report, subperiod review, sensitivity analysis
Robustness checks	Does the result survive reasonable changes in assumptions?	Pass, fail, conditional	Assumption grid and fragility notes
Portfolio construction	Does the signal survive the intended construction process?	Pass, fail, conditional	Construction tests, constraint analysis
Exposure control	Are beta, sector, factor, country where relevant, and concentration exposures understood?	Pass, fail, conditional	Risk report and exposure attribution
Sizing	Is the proposed size justified by evidence, uncertainty, liquidity, capacity and drawdown budget?	Pass, fail, conditional	Sizing memo and approval limit
Transaction costs	Are costs estimated consistently with turnover, horizon and execution style?	Pass, fail, conditional	Cost model, sensitivity test, implementation shortfall assumptions
Liquidity	Can trades be executed under normal and stressed assumptions?	Pass, fail, conditional, blocked	Liquidity report, participation assumptions
Capacity	Is the proposed capital level consistent with market depth and impact tolerance?	Pass, fail, conditional	Capacity range and assumptions
Execution assumptions	Does simulated execution match the intended trading workflow?	Pass, fail, conditional	Execution timing notes and operational workflow
Attribution readiness	Is attribution prepared by name, sector, factor, signal family, beta, concentration and cost?	Pass, fail, conditional	Attribution design and reporting format
Monitoring thresholds	Are performance, drawdown, turnover, cost, exposure, data and operations thresholds defined?	Pass, fail, conditional	Monitoring plan and escalation rules
Operational ownership	Are jobs, feeds, overrides, incidents and owner responsibilities defined?	Pass, fail, conditional	Operating procedure and incident log
Governance decision	Has the final decision been recorded with owner, date, version, limits and review schedule?	Approve, shadow, conditional, watchlist, block, reject	Approval memo and state record