

A TAXONOMY

# The Taxonomy of *Venture Conviction*.

*Seven archetypes, three vectors, one grammar for venture failure and conviction.  
A framework for anyone who reads a deck and has to decide what to do next.*

VECTOR I

## Solvency

*Physics decides.*

VECTOR II

## Structure

*Activity is not progress.*

VECTOR III

## Alpha

*The edge is in the observer.*

# Most failed deals share the same seven shapes.

*If you cannot name the shape of the deal in front of you, you are not pattern-matching. You are guessing in a vocabulary you have not bothered to learn.*

**T**he venture industry has industrialised due diligence, standardised term sheets, and automated portfolio monitoring. The one thing that still happens in tacit, unteachable, partner-by-partner whisper is the act of looking at a deal and deciding what kind of risk it carries.

That is the most expensive workflow in venture. It is also the most fixable.

Every deal that fails — and most do — fails in one of seven structural ways. They have names. They have signals. They have verdicts. Once you can name the shape, you stop arguing about whether a deal is good and start arguing about which mechanism is breaking.

This document is the grammar.

*Pattern-matching tells you a deal feels wrong. A taxonomy tells you which mechanism is breaking — and what to do about it.*

THE THESIS

# What it costs you not to have this framework.

---

*The absence of a shared structural vocabulary is not a stylistic preference. It is a quantifiable drag on attention, capital, and learning.*

## 01 Attention is your scarcest resource.

A typical fund sees 1,500–3,000 deals a year and writes 8–15 cheques. Without a structural pre-filter, every deal demands the same first-pass cognitive load. The taxonomy collapses 80% of inbound to a verdict in under sixty seconds.

## 03 Anti-portfolio learning becomes possible.

Most funds cannot say why they passed on the winners they passed on. Tacit reasoning leaves no trace. A structural classification creates an auditable record — and a fund that cannot audit its passes cannot improve them.

## 02 Disagreement becomes productive.

Two partners looking at the same deck will reach different verdicts for incompatible reasons. With a shared vocabulary they can argue about which archetype, not whether the deal is "good." That is the difference between investment-committee theatre and investment-committee work.

## 04 Junior judgement compounds faster.

Pattern-matching takes a decade to develop because it is built one deal at a time. Structural classification can be transmitted in a morning. The framework is not a substitute for experience; it is a scaffold that lets experience accumulate in a useful shape.

# Three vectors. Seven archetypes. Three verdicts.

*The entire framework on one page. If you remember nothing else, remember this map.*

## VECTOR I

### Fundamental Solvency

*Narrative cannot negotiate with physics.*

#### ARCHETYPE 1

The Hallucination

#### ARCHETYPE 2

The Fraud

IMMEDIATE PASS

## VECTOR II

### Structural Coherence

*Activity is not progress.*

#### ARCHETYPE 3

Strategic Incoherence

#### ARCHETYPE 4

The Subsidy Trap

#### ARCHETYPE 5

Legacy Debt

HIGH EXECUTION RISK

## VECTOR III

### Cognitive Alpha

*The edge is in the observer.*

#### ARCHETYPE 6

The False Negative

#### ARCHETYPE 7

The Deep Tech Winner

HIGH CONVICTION

# I

## Fundamental Solvency.

*Physics decides. Narrative does not negotiate.*

---

If unit economics do not close at any plausible scale, no amount of financing creativity fixes it. If technical claims are decoupled from technical capability, no amount of "trust the team" fixes it. Vector I is the floor. Most failed deals never get past it — and most rejected deals are rejected here for the right reasons.

ARCHETYPE 1

The Hallucination

Verdict · Pass

ARCHETYPE 2

The Fraud

Verdict · Pass

# The Hallucination: a solvency crisis dressed as a liquidity narrative.

*Each round solves a cash problem but never fixes the underlying economics. The deck shows growth; the unit economics tell a different story.*

**T**he company treats a fundamental unit-economics problem as a temporary cash-flow gap, raising successive rounds to mask deteriorating fundamentals. The narrative says *liquidity*; the arithmetic says *solvency*. This is the most expensive misdiagnosis in venture, because every additional round makes the next one larger and the eventual correction more violent.

The Hallucination does not look like a failed company. It looks like a company two rounds away from breakout — until you ask the one question its founders cannot answer.

**THE SIGNAL** Burn improves marginally with each raise and never reaches breakeven at any realistic customer count. Revenue growth and economic viability are conflated. The deck shows impressive top-line numbers next to a burn curve that accelerates, not decelerates, with scale.

**THE TELL** Founders defend the model by extending the timeline rather than tightening the unit. Profitability moves to the next round, then the round after that. Each financing extension is reframed as strategic patience.

**THE STRESS-TEST** Ask: *at what customer count do you reach contribution-margin positive?* If the answer requires a number that does not exist in the addressable market, this is a Hallucination. Walk.

## CASE · THE LAST-MILE ROLL-UP

### A nine-figure burn against the second law of thermodynamics.

A last-mile fleet roll-up burned through hundreds of millions chasing economies of scale that violated the physics of urban logistics.

The deck told a growth story. The unit economics told a heat-death story. Each round bought 9–14 months of runway and made the next round larger.

*It was never a financing problem. It was always a physics problem.*

*Narrative cannot negotiate with physics.*

# The Fraud: claims deliberately decoupled from capability.

*Not a team that failed to execute. A team that misrepresented their ability to execute from inception. The narrative is built to conceal, not reveal.*

The Fraud is structurally distinct from the well-meaning team that overpromised. Founders here construct a narrative engineered to defeat verification — not because verification is hard, but because the underlying capability does not exist.

The polish is the cover.

Sophisticated investors are not immune. They are sometimes more vulnerable, because they substitute social proof for technical audit.

**THE SIGNAL** Technical claims are impossible to verify in the pitch context. Demos are staged. Data is selectively presented. Third-party validation is cited but cannot be independently confirmed. References lead to the company's network, not to disinterested operators.

**THE TELL** Founders avoid specificity under pressure. Questions about reproducibility, regulatory approval, or live customer verification produce *narrative pivots*, not evidence. Deflection is fluent and rehearsed.

**THE STRESS-TEST** The Fraud often has the highest Presentation Score in the pipeline. Polish is structurally indistinguishable from substance unless you demand verification. So demand verification — read the report, not the reference.

## CASE · THERANOS

### When narrative construction outruns technical feasibility.

Theranos is the canonical study. The technology did not exist; the narrative did. Sophisticated investors, board members with national security clearance, and Walgreens all signed up.

The lesson is not that fraud is rare. The lesson is that polish hides everything until you ask for the report rather than the reference.

*The single question: "May we read the validation report?" If the answer is anything other than yes, walk.*

*Polish is structurally indistinguishable from substance — unless you demand verification.*

# II

## Structural Coherence.

*Activity is not progress.*

---

The economics may close, the team may be honest, the market may be real — but the company is set up in a way that makes execution structurally unlikely. Vector II is the murkiest layer to diagnose because the deck looks competent. The work is to ask whether the architecture matches the capital, and whether the *incentives match the architecture*.

### ARCHETYPE 3

#### Strategic Incoherence

Verdict · High Execution Risk

### ARCHETYPE 4

#### The Subsidy Trap

Verdict · High Execution Risk

### ARCHETYPE 5

#### Legacy Debt

Verdict · High Execution Risk

# Strategic Incoherence: a multi-front war on seed capital.

*The company tries to build platform infrastructure, marketplace liquidity, proprietary technology, and brand simultaneously — each of which alone would need a Series B.*

**T**he deck lists three or more simultaneous strategic priorities, each with its own team, timeline, and capital requirement. The Gantt chart shows parallel tracks that cannot actually be staffed by the headcount the raise will fund. Complexity is sold as moat. Focus trade-offs are not acknowledged because acknowledging them threatens the fundraising narrative.

The right rebuttal is not *your vision is too small*. It is *your vision is right; your sequencing is wrong*.

**THE SIGNAL** Three or more concurrent priorities each requiring Series B-level capital. Headcount plan does not match the timeline plan. Each leg of the strategy depends on the others working — but no leg is funded to standalone viability.

**THE TELL** Founders defend complexity as moat. *"We need all of this to win"* is the classic line. Asked to sequence, they refuse — sequencing means admitting that something must wait.

**THE STRESS-TEST** Ask: *"What is the one thing that must work?"* If the founder cannot answer, complexity is not their moat. It is their anchor.

## ALSO KNOWN AS · THE SUPER-APP DELUSION

**Pre-seed startups launching three to five distinct business lines simultaneously.**

Each business line requires Series B-level capital and organisational focus to execute. Combined into a single pre-seed raise, they form an architecture that cannot be staffed and cannot be sequenced.

The Super-App Delusion is most common in geographies where one or two genuine super-apps have succeeded — survivorship bias rationalising the strategy.

*The fix is not less ambition. It is sequencing.*

**Complexity is not their moat. It is their anchor.**

# The Subsidy Trap: government grants substituting for market validation.

*Growth is a function of regulatory capture and grant extraction, not organic demand — building a business that cannot survive contact with unsubsidised competition.*

Revenue is primarily grants, government contracts, or regulatory mandates. Commercial customers are scarce or pilot-stage. The growth chart shows a funding timeline, not a sales pipeline.

The subsidy accelerates early traction and creates the appearance of product-market fit — then creates a dependency that venture capital cannot fix, only deepen. Watch for any deck where the customer logos are ministry seals.

**THE SIGNAL** Revenue mix dominated by grants, government contracts, or regulatory mandates. Commercial pipeline is thin or pilot-stage. The growth chart aligns with grant cycles, not sales cycles.

**THE TELL** Strip the subsidy and ask: *who pays, at what price, and why now?* If the answer is uncertain, the company has optimised for grant-writing, not market creation.

**THE STRESS-TEST** Ask: *"What are your fully-loaded unit economics without the grant?"* If the answer does not exist, this is structural dependency, not strategic advantage.

## WHERE IT APPEARS

Climate. Defence. Govtech. Wherever subsidy precedes market.

Climate tech, defence tech, govtech — all sectors where grants and contracts can substitute for early commercial traction. The subsidy is real; so is the dependency.

Government validation is not market validation. The customer who can be replaced by the next administration is not a customer; they are a counterparty.

*The fix is to underwrite the post-subsidy business. If it does not stand without the grant, the grant is the business.*

*Government validation is not market validation.*

# Legacy Debt: digital transformation slowed by internal cannibalisation.

*The innovation agenda conflicts with the incentives, capabilities, and political economy of the existing business — creating organisational antibodies that reject the transformation.*

The innovation unit has a budget, a team, and a roadmap. But it reports to a business unit leader whose P&L depends on the existing product. The innovation team is structurally subordinated to the thing it is meant to replace. The transformation has been captured before it started.

This is the most governance-sensitive archetype. It requires board-level structural intervention, not better product management.

**THE SIGNAL** Innovation reports through a leader whose compensation is tied to existing revenue. Budget approval still sits with the parent business. Cap-table independence is cosmetic; operational dependence is total.

**THE TELL** Ask who has budget authority over the transformation. If the answer is a person whose compensation is tied to the existing revenue stream, the transformation has been captured.

**THE STRESS-TEST** Diagnose by following *reporting lines*, not org charts. Exceptional governance can overcome Legacy Debt — but it must be acknowledged at the board level, not managed around at the operational one.

## WHERE IT APPEARS


Corporate venture. Series B-plus transformations. Spinouts that did not actually spin out.

Most common in corporate venture and Series B-plus digital transformation plays where the incumbent's existing customer base is simultaneously the asset and the constraint.

Also frequent in spinouts whose cap table looks independent but where budget approval still routes through the parent.

*The diagnostic is reporting lines, not org charts.*

*The transformation has been captured before it started.*



# Cognitive Alpha.

*The edge is in the observer.*

---

By the time a deal reaches Vector III, physics and structure have cleared. The remaining question is whether you can see something the market cannot. Two archetypes live here: deals the consensus has mispriced, and deals where technical difficulty itself is the moat. Both reward concentration. Both punish hedging. *If you have correctly identified the edge, a small position is irrational.*

ARCHETYPE 6

The False Negative

Verdict · High Conviction

ARCHETYPE 7

The Deep Tech Winner

Verdict · High Conviction

# The False Negative: rejected by consensus, validated by structural judgment.

*The market systematically misprices the opportunity due to category confusion, timing mismatch, or incomplete mental models. The contrarian sees what the generalist cannot.*

**T**he deal has been passed by multiple credible investors for consistent reasons. The reasons are not *team* or *market size* — they are category-level objections that reveal a mental-model gap, not a business-model gap. The market is looking at the deal through the wrong lens.

The investor who sees the False Negative can articulate exactly what the market is getting wrong, and why. The alpha is not a hunch; it is structural judgment from domain knowledge the rejecting investors do not have.

**THE SIGNAL** Multiple credible investors have passed for consistent reasons. The pattern of objection is category-level, not company-level. Investors are answering a different question than the one the company is asking.

**THE TELL** You can articulate the consensus error precisely — name the category, the mispricing, the missing mental model. If you cannot, you do not have an edge; you have an opinion.

**THE STRESS-TEST** False Negatives reward portfolio concentration, not diversification. If you have correctly identified the consensus error, holding a small position is irrational. *The risk is in the judgment, not the allocation.*

## CASE · AIRBNB, 2008

The investors who said yes were not braver. They had a different model of what the company was.

Rejected by most institutional investors as housing-crisis arbitrage rather than a platform-enabled trust network. The structural judgment was in the trust layer, not the room.

The consensus saw a real-estate side-hustle; the contrarians saw the formation of a trust marketplace that had no precedent in the category they were comparing it to.

*The category was wrong, not the deal.*

*The risk is in the judgment, not the allocation.*

# The Deep Tech Winner: difficulty as moat, paradigm as wedge.

*Technical barriers limit competition. Market timing creates a narrow window. Business model innovation compounds the technical advantage. All three at once is the rarest signal in deal flow.*

Technical barriers are real and independently verifiable. The window of opportunity is narrow because it depends on a recent paradigm shift that incumbents cannot respond to quickly. The business model amplifies, rather than dilutes, the technical advantage.

Evaluating a Deep Tech Winner requires domain expertise most generalist investors do not have. Without that capability, this is speculation in deep-tech costume.

**THE SIGNAL** Real, independently verifiable technical barriers. A paradigm shift opening a window incumbents cannot close in time. Business model that compounds — not dilutes — the technical edge.

**THE TELL** The alpha derives from *your ability to make the technical-feasibility judgement*. If your team cannot make that judgement, the deal is not in your circle of competence — regardless of how attractive the narrative.

**THE STRESS-TEST** Deep Tech Winners justify portfolio concentration and patient capital. The uncertainty phase is not a risk to be managed — it is the source of the return. Impatient capital de-risks the opportunity by pressuring premature commercialisation.

## RARITY · THE HARDEST DIAGNOSIS

Roughly two of every hundred-and-thirty cases.

Genuine technical novelty combined with business-model innovation is the rarest signal in deal flow.

If a portfolio claims more than one or two Deep Tech Winners per vintage, the diagnosis is almost certainly wrong. Most candidates are Hallucinations in deep-tech costume.

*The rarity is the signal. The frequency is the trap.*

*Most candidates are Hallucinations in deep-tech costume.*

## SYNTHESIS

# The seven archetypes at a glance.

Every deal narrative maps to one dominant archetype. Most map to more than one.

ARCHETYPE	VECTOR	CORE RISK · SIGNAL	VERDICT
1 The Hallucination	I · Solvency	Burn accelerates with scale	PASS
2 The Fraud	I · Solvency	High presentation, low logic	PASS
3 Strategic Incoherence	II · Structure	More than three pre-seed pillars	EXEC RISK
4 The Subsidy Trap	II · Structure	Revenue is grants, not customers	EXEC RISK
5 Legacy Debt	II · Structure	Innovation subordinated to legacy P&L	EXEC RISK
6 The False Negative	III · Alpha	Consensus rejection + mental-model gap	HIGH CONV.
7 The Deep Tech Winner	III · Alpha	Paradigm shift + narrow window	HIGH CONV.

Five of seven archetypes produce a stop verdict. That is not a flaw of the framework — it is its primary operational value. Most deal flow is bad. **Protecting attention is the first job.**

# When archetypes overlap, each one makes the others harder to escape.

*Most deals do not have one problem. They have three — and the combinations are where capital actually dies.*

COMPOUNDING · 1+4

## The Hallucination + The Subsidy Trap

*Broken unit economics masked by grant revenue.*

The grant creates the illusion of product-market fit while the underlying economics deteriorate. Investors mistake regulatory capture for commercial validation. By the time the subsidy ends, three vintages of capital are already stranded.

COMPOUNDING · 2+3

## The Fraud + Strategic Incoherence

*Complexity used to obscure verification.*

Multi-front strategy makes technical claims impossible to verify because nothing is far enough along to audit. The complexity is the cover, not the strategy. Each individual claim is unverifiable; the combination is unfalsifiable.

MAX CONVICTION ·  
6+7

## The False Negative + The Deep Tech Winner

*The rarest and highest-conviction pattern.*

A deal the market has passed on for category-confusion reasons, where the technical barriers are real and the window is narrow. These are the Airbnb-level opportunities. They are also the most often misdiagnosed: *one is real for every dozen claimed.*

# The three-question filter, before any detailed diligence.

Three questions, asked in order. Each one routes the deal to a vector. Each vector dictates the response.

**Q1 · SOLVENCY** Does the business model survive contact with physics and economics?

✓ Unit economics reach contribution-margin positive at a realistic customer count. No assumption requires a market that cannot exist.

✗ Burn accelerates with scale. Claims are technically unverifiable.

→ Vector I

**PASS**

**Q2 · STRUCTURE**

Does the strategic architecture match the available capital and team?

✓ One clear priority. Headcount plan coherent with the raise. Revenue validated by paying customers, not grants.

✗ Multiple simultaneous priorities. Grant-dependent revenue. Transformation conflicts with existing business incentives.

→ Vector II

**HIGH EXECUTION RISK**

**Q3 · ALPHA**

Do you possess structural judgment the market does not?

✓ You can articulate exactly what the consensus is getting wrong, and why. The technical-feasibility judgement is within your domain competence.

✗ You have a hunch. You like the team. The market is large.

→ Vector III

**HIGH CONVICTION**

# Three narratives. Three diagnoses. Try before reading the answers.

Cover the right column. Read each narrative. Name the vector and the archetype. Then check.

**A** AgriStack raises \$8M Series A. Revenue is 90% government grants from three ministries. Two paying commercial customers, both at pilot pricing. Founder argues grant validation proves product-market fit.

**Vector II · Archetype 4 · Subsidy Trap**

Government validation is not market validation. The pilots are the only commercial signal — and they are not at full price. Strip the grants and ask who pays.

**B** MediScan AI claims proprietary diagnostic technology validated by "leading hospitals." Demos are pre-recorded. Independent testing requests are redirected to an NDA. Presentation Score: exceptional.

**Vector I · Archetype 2 · The Fraud**

Polish without verification. The single question: "May we read the validation report?" If the answer is anything other than yes, walk.

**C** Twelve top-tier funds passed on this logistics platform in 2021 for "thin margins in a commodity market." Two firms with deep supply-chain expertise saw a trust-layer network effect the generalists missed. The company is now #2 in the category.

**Vector III · Archetype 6 · The False Negative**

The market priced the deal as logistics. The contrarians priced it as trust infrastructure. The alpha was in the category, not the cash flows.

## KEY TAKEAWAYS

# Five things to remember after the deck is closed.

*If you forward only one page from this document, this is the page to forward.*

- 01** Startup failure follows a universal grammar. Seven archetypes, three vectors. Pattern matching is not enough — structural classification is.
- 02** Five of seven archetypes produce a stop verdict. Most deal flow lives in Vectors I and II. Protecting attention is the primary operational value of the framework.
- 03** The two Alpha archetypes reward concentration, not diversification. If the consensus error is correctly identified, a small position is irrational.
- 04** Compound risk is the hardest pattern to detect — and the most expensive to miss. When archetypes overlap, each one makes the others harder to escape.
- 05** A taxonomy is not a substitute for judgement. It is the scaffold that makes judgement teachable, transmissible, and auditable.

# Information is cheap. Judgement is the last scarce asset.

*The taxonomy is one piece of a larger argument: that the venture industry has industrialised everything except the act of deciding.*

Due diligence has been industrialised. Term sheets have been standardised. Portfolio monitoring has been automated. The core function — transforming information into defensible investment decisions — remains trapped in tacit knowledge that walks out the door when senior partners retire.

A taxonomy of conviction is the start. It makes one part of the judgement act explicit, shareable, and auditable.

The next pieces — systematic stress-testing of brittle assumptions, anti-portfolio learning loops, scaling senior judgement across an investment team — are downstream of this same shift.

## A NOTE ON ENFORCEMENT.

This document describes the taxonomy. It does not describe the engine that enforces it. The same seven archetypes that take a partner ten years of pattern-matching to internalise are the ruleset askOdin's RUNE compiler executes deterministically against a data room — producing a Clarity Score and an audit trail in minutes rather than meetings. The taxonomy is the law. RUNE is the enforcement.

*The infrastructure is not for finding more deals. It is for making the deals you already see legible to the people who have to act on them.*

# Terms used in this document.

---

*A short glossary, in service of precision rather than gatekeeping.*

---

## Anti-portfolio

The deals a fund passed on that subsequently became material winners. The structural reason for each pass — when recorded — is the highest-quality training data a fund owns.

---

## Contribution margin

Revenue per unit minus the variable cost of producing that unit. A business reaches contribution-margin positive when each new customer pays for the cost of serving them.

---

## Paradigm shift

A discontinuity that changes which capabilities are valuable. Recent paradigm shifts open windows incumbents cannot close in time — the precondition for Archetype 7.

---

## Clarity Score

A deterministic measure of how legible a deal's underlying claims, capital structure, and unit economics are when stress-tested. The structural counterpart to Presentation Score.

---

## Stress-test

The single question that, asked sharply, distinguishes one archetype from another. Each archetype in this document has one.

---

## Burn rate

The pace at which a company consumes cash. Improves with scale only if unit economics close. If burn accelerates with scale, the model is broken.

---

## Mental-model gap

A consensus error rooted not in missing data but in the wrong category. The contrarian and the consensus see the same facts but classify them differently.

---

## Presentation Score

A measure of how polished a pitch is. Empirically uncorrelated with business quality. Becomes diagnostic only when compared against the Clarity Score.

---

## Narrative Masking

The diagnostic delta between Presentation Score and Clarity Score. A high Presentation Score against a low Clarity Score is the signature of Archetype 2 — The Fraud.

---

## Unit economics

The profit-and-loss of a single transaction or customer relationship. The first thing physics decides for a business.

# Seven archetypes. One grammar.

*Judgement is the last unscalable asset. The first step in scaling it is naming it.*

## FOR INVESTORS

Use the three-question filter on your next inbound deck before assigning diligence.

## FOR ANALYSTS

Classify every deal you read this week. Compare your verdicts with the partner's. Note the gap.

## FOR ALLOCATORS

Ask any GP to name their last ten passes by archetype. The answer is a portrait of their judgement.