

LANTERN LABS

INVESTMENT RESEARCH

www.lanternlabs.net

Finding the Real Ones

Decision Framework Memo

A Startup Evaluation Framework for the AI Age
July 2026

A Decision Framework for Separating Durable AI-Era Companies from Disposable AI Wrappers

Structured thinking study — original analysis. Not investment advice. Illustrative cases only.

Decision Intelligence, not Decision Making — this report informs the client's decision; it is not a recommendation.

Prepared by Lantern Labs

Contents

EXECUTIVE BRIEF	3
Decision Spine — The Report in 60 Seconds	4
Ground Truth — The Evidence Base	6
Why the Old Filter Breaks	7
02 • The Six-Part Formula	8
The Elements	8
Old Tradeoff	8
New Capability	8
Unresolved Truth-Conflict	8
Right Timing	9
Hard Constraint	9
Nested Conflict Depth	9
03 • Historical Proof: The Thesis Before AI	11
The 2036 Hindsight Test — Obvious Later, Non-Obvious Now	12
THE PRACTICAL FILTER	13
05 • The Eight Questions	13
06 • What Fails the Framework	14
07 • Worked Examples	16
Example A — AI Education (“The Formation Problem”)	16
Example B — AI Healthcare Administration (“The Documentation Trap”)	17
08 • Hidden Principles	17
Principle 1 — Capability Abundance Shifts Value to Judgment	17
Principle 2 — Constraints Are Not the Enemy of Invention	18
Principle 3 — The Wedge Predicts the Platform Only If the Conflict Chain Is Deep	18
Case Study — Framework Applied: PASS / WATCH / KILL	18
The Idea: AI Prior Authorization and Clinical Documentation for Specialty Care	18
Six-Part Formula Applied	19
Buyer and Commercial Path	20
Failure Pattern Check	21
SYNTHESIS	23
Anti-Thesis — The Strongest Case Against This Framework	23
09 • The Synthesis	23
What Changes By Audience	23
The Remaining Truth	23
Closing	24
CHECKLIST — FINDING THE REAL ONES	25
Part A — The Six Elements (structural test)	25
Part B — The Failure Pattern Check	25
Part C — The Buyer Test (the implicit ninth question)	26
Sources	26
Disclaimer	26

EXECUTIVE BRIEF

The conclusion, the reasoning, and the one thing to act on

CORE THESIS

In the AI age, execution is cheap. Judgment is scarce. That divergence — not AI capability itself — is the defining strategic variable for anyone building or backing a startup today.

THE OLD FORMULA

Most startup evaluation stops here:

Good team + large market + working product = fundable company.

This is not wrong. It is insufficient. It cannot distinguish between the 10,000 AI wrappers launched in 2025 and the handful of companies that will compound over decades.

THE NEW FORMULA

Old tradeoff + new capability + unresolved truth-conflict + right timing + hard constraint + nested conflict depth = serious startup opportunity.

Each element does specific work. Remove one and you get a product. Keep all six and you get a company.

THE SHARPEST MECHANISM

The market conflates “AI makes this possible” with “this is a real opportunity.” These are different claims. AI makes thousands of things possible simultaneously. The question is which possibilities sit inside a structure that is: timely (why now, not two years ago), defensible (hard enough to protect), and expandable (the first problem opens into a larger system). The six-part formula is a filter for exactly that structure.

WHY NOW

The 100× cost reduction in AI reasoning since January 2023 has already happened. Every idea filter built before 2024 was calibrated for a world where execution was the bottleneck. That world no longer exists. When execution is cheap, the bottleneck shifts to judgment — the capacity to identify which things are worth building, for whom, under what constraints, at what moment. The founder who can answer that precisely has a structural advantage that cannot be replicated by better tooling.

The old bottleneck	The new bottleneck
Can we build it?	Should we build it — in this form, for these people, right now?
Access to engineers	Access to founders with sound judgment
Capital efficiency	Judgment efficiency
Time to first product	Time to right problem

ONE CAVEAT BEFORE THE FORMULA

The six-part formula evaluates opportunity structure. It does not replace buyer urgency. A structurally elegant idea is not a market until a specific buyer feels the conflict acutely enough to pay, switch, or embed the solution. The framework identifies whether an opportunity deserves serious attention. It does not tell you whether the market is ready to act on it today.

THE NON-OBVIOUS MOVE — BY AUDIENCE

You are	The non-obvious move
Founder	Stop evaluating ideas by whether AI makes them buildable. Evaluate them by whether the opportunity structure — tradeoff, timing, constraint, nested conflicts — is genuinely present. Most ideas that feel fundable fail this test.
Investor	The mispriced category is constrained domains: healthcare, education, legal, infrastructure, regulated industries. Easy AI is already crowded and cheap. Hard AI, where the constraint is the moat, is underinvested.
Operator / Executive	Apply the “nested conflict” question to your own category. After the problem your company solves today — what more important problem becomes visible? That answer is your strategic roadmap and your acquisition thesis simultaneously.

Decision Spine — The Report in 60 Seconds

Item	Answer
Core thesis	The bottleneck has shifted from buildability to worth-building. In the AI age, nearly anything can be built. The scarce resource is the judgment to know what deserves to be.
The formula	Old tradeoff + new capability + truth-conflict + right timing + hard constraint + nested conflict depth = serious startup opportunity.
What passes	Constrained, timely, expandable ideas where AI breaks a real tradeoff without violating a genuine truth-conflict — and where the first problem opens into a larger system.
What fails	Thin AI wrappers (no constraint), cool demos (no buyer urgency), broad platforms from day one (no wedge), regulated-market fantasies (constraint present, no credible path through it), timing-dependent ideas entered too early.
Mispriced category	Hard AI in constrained domains: healthcare, education, legal, regulated industries, physical-world automation. Easy AI is already crowded. Hard AI, where the constraint is the moat, is underinvested.
2036 watchlist	Verification infrastructure, AI-native education, healthcare workflow systems, governance and compliance rails, judgment-training platforms, constrained robotics, science research operating systems.

Item	Answer
Primary risk to the framework	Some category-defining companies begin with distribution advantage, founder-market pull, or execution speed — before the full structure is visible. The formula identifies good structure. It does not replace customer discovery, timing luck, or execution quality.
How to use this report	Executive Brief for the thesis. Historical Proof for pattern evidence. The Practical Filter for the tools. The Checklist to evaluate a specific idea.

Ground Truth — The Evidence Base

The core thesis — execution is cheap, judgment is scarce — is not assertion. It rests on seven verifiable facts. Each one changes what the right startup filter looks like.

Claim	What it means	Ref
GPT-4 class reasoning cost fell approximately 100× between January 2023 and early 2025	Intelligence is becoming a commodity faster than any pre-2024 strategy anticipated. Every cost-dependent moat built on “we have AI access” has already eroded.	[1]
AI startup formation grew faster than any prior technology wave — with the majority of new AI products being thin wrappers on foundation models	When execution is cheap, the field floods with undifferentiated products. Most will not survive first contact with real enterprise procurement.	[2]
Enterprise AI adoption is rising but most implementations stall at pilot stage — the constraint is workflow integration, not model capability	Capability is not the bottleneck. Accountability, change management, and workflow trust are. The structural shift: from model-bottleneck to organisation-bottleneck.	[3]
EU AI Act major enforcement begins August 2, 2026 — with high-risk category rules covering healthcare, education, and critical infrastructure	Regulation is no longer theoretical. Compliance cost and procurement friction are now structural conditions — exactly the kind of hard constraint that becomes a moat for the first credible operator.	[4]
AI investment is shifting toward the physical economy — energy, robotics, logistics, and agriculture received the fastest-growing share of AI capital in 2025–2026	Capital has already reached the same structural conclusion the framework predicts: easy software AI is saturated; constrained physical domains are the next frontier.	[5]
Documented AI incidents — hallucinations, bias, safety failures — are rising in frequency even as model capability improves	As AI output volume increases, the capacity to evaluate what it generates becomes more critical, not less. Verification infrastructure is a structural response to a structural problem.	[6]
Healthcare administrative burden costs US health systems an estimated \$265B annually — the largest operational cost category outside direct patient care	The AI opportunity in healthcare is administrative relief, care coordination, and documentation. The buyer, budget, and regulatory pathway are all defined.	[7]

One-line summary of what the evidence base establishes:

The opportunity is not AI in general. It is AI deployed inside a hard constraint — regulatory, trust, workflow, or physical — where the constraint protects the operator who can navigate it and prices out the competitor who cannot.

Why the Old Filter Breaks

The standard startup evaluation checklist was built for a world where execution was the bottleneck. In that world, the key questions were:

- Is the team capable of building this?
- Is the market large enough to justify the effort?
- Is there evidence that customers want it?

These remain necessary. They are no longer sufficient.

Why they break in the AI age:

Old question	Why it fails now
“Can this team build it?”	Nearly any team can build a working MVP using AI tools. Build capability is no longer a differentiator.
“Is the market large enough?”	TAM estimates are almost always too large. The relevant question is not the total market but the initial wedge — specific, concrete, and winnable.
“Do customers want it?”	Customers say they want hundreds of things they will never pay for, switch for, or embed in their workflows. Desire is not adoption.
“Is there a moat?”	“AI moat” is the most common false claim in pitches today. A feature is not a moat. A distribution advantage is not a moat. Moats in AI come from data, trust, workflow integration, and regulatory relationships — not from model access.

The deeper problem is that these questions evaluate whether a company can exist. They do not evaluate whether the underlying opportunity structure makes it worth building.

The opportunity structure question is different:

Why does this tradeoff exist? What changed that makes it breakable now? Why haven't the incumbents already broken it? What makes this hard enough that success cannot be easily copied? And — most importantly — does solving the first problem open into something larger?

These are the questions the six-part formula is designed to surface.

Insight

The average AI founder evaluates ideas by asking whether they can be built. The best AI founders evaluate ideas by asking whether the structure beneath them is real — and whether the timing, the constraint, and the depth of conflict make them worth building now, in this form, for these people.

02 • The Six-Part Formula

The Elements

Old Tradeoff

Every durable business exists because the world has accepted a painful compromise — two things people want simultaneously that the current market cannot provide together. The tradeoff is not the same as the problem. The problem is a symptom. The tradeoff is the structural reason the symptom persists. The right question is not “what do people complain about?” but “what painful compromise have people silently accepted as permanent?”

Weak version	Strong version
“Healthcare is expensive”	“Doctors must choose between seeing more patients or knowing each patient well — the time constraint forces the compromise”
“Education isn’t personalized”	“Personalised learning requires a 1:1 human tutor most families cannot afford — so schools optimise for the median student”
“Freight is unreliable”	“Speed and cost trade off structurally in logistics because the incentive systems of carriers and shippers point in opposite directions”

New Capability

The new capability is the specific development — technological, regulatory, economic, behavioral — that makes the tradeoff now breakable in a way it was not before. This is not always AI: a cost drop, regulatory change, behavioral shift, or infrastructure change all qualify. The capability question must be answered precisely. “AI is better now” is not an answer. “Inference costs dropped 100× since January 2023, making continuous AI-assisted tutoring economically viable for the first time” is an answer.

Unresolved Truth-Conflict

The truth-conflict is the deepest element — and the most commonly skipped. It is not a simple problem but a tension between two things that are independently true: solving one, if done naively, violates the other. The market stays stuck because both truths are real and neither can be abandoned. The opportunity is to build something that allows both to survive simultaneously — not a compromise, but a genuine synthesis.

Domain	Truth A	Truth B
Education	Students need personalized help and feedback	Students also need shared standards, discipline, and human formation
Healthcare	Patients need more access to medical expertise	Medical expertise requires trust, accountability, and human judgment
Legal	Clients need fast, affordable legal guidance	Legal advice requires accountability, context, and licensed judgment

Insight

The strongest startup ideas are not solutions to simple problems. They are syntheses of genuine truth-conflicts. The synthesis is the product. The fact that both truths are real — and that naive solutions violate one of them — is what creates durable differentiation.


Right Timing

A structurally breakable tradeoff is not the same as the right moment to break it. The strongest “why now” answers name a precise recent change — not a general trend — that opens the window. “AI is improving” is not a timing answer. “The EU AI Act enforcement deadline in August 2026 creates immediate compliance pressure that makes governance tools mandatory, not optional” is. The test: *why could this not have been built and scaled two years ago?* If the answer is weak, the window is not yet open — or has already closed.

Hard Constraint

The constraint is what makes the opportunity defensible. Easy AI opportunities — those with no meaningful constraints — attract thousands of competitors and compress to zero margin within 18 months. The best AI opportunities sit where high AI capability meets high real-world constraint simultaneously. The question is not “does a constraint exist?” but “is this constraint hard enough that a well-resourced competitor cannot replicate what we build in 12 months?” One critical distinction: a constraint becomes a moat only when the startup has a credible path to operate inside it better than incumbents. Otherwise it protects the incumbents, not the new entrant — the regulated-market fantasy.

Type of constraint	What it protects	Example domain
Regulatory	Access, liability, compliance relationships	Healthcare, financial advice, legal, education
Trust	Long-term relationships, track records, certification	Medical, legal, financial, B2B enterprise
Workflow integration	Deep embedding in existing processes	EHR systems, ERP software, construction management
Domain expertise	Irreplaceable human judgment layered on AI	Investment research, scientific research, high-stakes legal




Nested Conflict Depth

The final and most predictive element. Small companies solve one conflict. Great companies discover that solving the first reveals a deeper one — and that one reveals a still-deeper one. The test: *after we solve the first problem, what more important problem becomes visible?* A shallow idea runs out of depth. A platform idea finds the first problem was the wedge into a much larger system.

Company	First conflict	Second conflict	Third conflict
Uber	Riders need rides; taxis unreliable	Riders need trust in strangers	Drivers need earnings stability; cities need safety
Airbnb	Travelers want affordable lodging; hosts have spare rooms	Guests need trust in strangers' homes	Hosts need protection; cities need housing regulation
Stripe	Developers need to accept payments	Businesses need financial infrastructure	Platforms need a complete financial stack

The question at evaluation is not whether the company will face these conflicts. They will. The question is whether solving the first conflict opens into a large enough system that the subsequent conflicts support a platform — not just a feature.



03 • Historical Proof: The Thesis Before AI

The six-part formula is not AI-specific. AI makes it more urgent, but the pattern appears across every company that defined its category. Five examples, different eras, same formula.

Example	Old tradeoff	New capability	Truth-conflict	Hard constraint / moat
Ford Model T	Cars were desirable but unaffordable — artisanal production meant only the wealthy could own one	Assembly-line manufacturing and process standardization	People wanted personal mobility AND mass-market affordability	Manufacturing system, labor organization, supplier coordination, dealer network
iPhone	Phones were portable but limited; computers were powerful but not with you	Touchscreen, mobile chips, app ecosystem, hardware-software integration	People wanted full computing power AND simplicity, beauty, and mobility	Hardware integration, software platform, carrier relationships, developer ecosystem, supply chain
Amazon	Physical retail had trust and immediacy; catalog shopping had selection but poor discovery and slow delivery	Internet commerce, search, logistics software, cloud infrastructure	Customers wanted infinite selection AND trust, low prices, and reliable delivery	Fulfillment, supplier relationships, customer trust, logistics, data infrastructure
Airbnb	Travelers wanted affordable, local lodging; hosts had unused space — but strangers could not trust each other	Online marketplaces, digital payments, reputation systems, identity verification, mobile booking	Travelers wanted cheaper authentic stays AND safety, reliability, and trust	Trust infrastructure, host supply, local regulation, insurance, marketplace liquidity
Stripe	Businesses needed to accept online payments, but payment infrastructure was complex, slow, and developer-hostile	APIs, cloud software, developer-first infrastructure, digital commerce growth	Businesses wanted payments to be simple AND they required compliance, fraud prevention, and banking reliability	Compliance, fraud, banking partnerships, uptime, global payment rules, developer trust

The shared lesson across all five:

These companies were not merely better products. They were syntheses of genuine contradictions — made possible by a new capability, at the right moment, inside constraints hard enough to protect them. Each one entered a contested space where two true things were in tension, built the synthesis, and discovered that solving the first conflict revealed a deeper one worth more than the first.

Insight

The formula is not merely descriptive after the fact. It is diagnostic before the fact. The question is not whether an idea sounds big. The question is whether the structure beneath it resembles the structure that has produced category-defining companies before: a breakable tradeoff, a new capability, a genuine truth-conflict, precise timing, hard constraints, and a conflict chain deep enough to become a platform.

The 2036 Hindsight Test — Obvious Later, Non-Obvious Now

The most valuable opportunities rarely look obvious when they first appear — too early, too constrained, too narrow. Ten years later, they look inevitable. Imagine standing in 2036: which categories were hiding in plain sight in 2026, dismissed as too constrained for venture capital, precisely because the constraint was the moat?

2036 obvious category	The old tradeoff	The truth-conflict	The hard constraint	The first wedge
AI-native education	Personalization vs affordability	Personalized learning vs human formation	Child safety, trust, curriculum standards	AI tutor for homeschool math — one subject, one segment
Verification infrastructure	Output speed vs confidence in truth	AI generation abundance vs reliable knowledge	Auditability, expert review, provenance	Verification layer for AI-generated professional work
Healthcare workflow OS	Patient volume vs care depth	Automation vs human accountability	HIPAA, EHR integration, malpractice	Specialty documentation tool for one clinical workflow
AI governance and compliance rails	Deployment speed vs institutional accountability	AI leverage vs legal/ethical responsibility	Regulation, audit trails, insurance, liability	Model-risk monitoring for one regulated industry
Professional judgment training	AI automates junior work vs firms need a future expertise pipeline	Productivity vs apprenticeship	Credentialing, simulation, assessment, liability	Decision simulator for one profession (law, finance, medicine)
Constrained physical robotics	Automation is valuable vs brittle and expensive	Physical autonomy vs safety and reliability	Hardware, deployment, maintenance, customer trust	One task in one constrained environment (warehouse, agriculture, lab)
Science research OS	Scientific speed vs truth verified by reality	Hypothesis abundance vs experimental validation	Wet labs, reproducibility, domain expertise, peer review	Literature-to-experiment workflow for one scientific domain

What the pattern teaches:

The obvious categories of 2036 are not the easiest AI applications. They are the ones where AI capability collides with a high-stakes human constraint — formation, trust, accountability, safety, expertise, privacy. The constraint is why each opportunity looks unattractive early. It is also precisely why each becomes defensible later.

Insight

The 2036 hindsight test changes the question. It does not ask which AI products look exciting today. It asks which unresolved tensions will look irresponsible to have ignored ten years from now. Easy AI becomes crowded. Hard AI becomes infrastructure. The founder's job is to find the intersection where those two facts meet.

THE PRACTICAL FILTER

05 • The Eight Questions

#	Question	What a weak answer looks like	What a strong answer looks like
1	Old tradeoff: What painful compromise have people silently accepted?	“Healthcare is too expensive”	“Doctors must trade patient volume for patient depth — the time constraint is structural”
2	New capability: What changed that makes this tradeoff breakable now?	“AI is getting better”	“Inference costs dropped 100× since 2023, making continuous AI assistance economically viable at scale for the first time”
3	Truth-conflict: What two independently true things are in tension?	“People want X but can’t afford it”	“Children need personalized help AND shared formation; naive AI tutors solve the first but violate the second”
4	Synthesis: What product lets both truths survive simultaneously?	“An app that does both”	“AI handles adaptive practice; humans handle formation, mentorship, and shared standards — a deliberate division of labor”
5	Timing: Why now — specifically, not generally?	“The market is ready”	“Remote work normalized async learning; AI quality crossed the tutoring threshold in Q1 2025; homeschooling is at a 20-year high”
6	Constraint: What makes this hard enough that success cannot be copied easily?	“Our tech is unique”	“HIPAA compliance, EHR integration, clinical workflow trust, and malpractice accountability — four compounding barriers, none of which are API calls”
7	Nested conflict: What deeper conflict appears after the first one is solved?	“We’ll cross that bridge when we get there”	“After personalized tutoring: how does this form character, not just knowledge? That’s the platform question — and it’s bigger than the wedge”

#	Question	What a weak answer looks like	What a strong answer looks like
8	Wedge: What is the smallest concrete version that proves the thesis?	“We’ll start with education broadly”	“AI mastery tutor for homeschool families in K-8 math, with parent dashboard and human-designed curriculum — one subject, one segment, one proof point”

Insight

The eight questions are a filter, not a scorecard. An idea that answers seven of eight well is a stronger opportunity than one that answers all eight weakly. The goal is not to check every box — it is to stress-test the underlying structure before committing resources to it.

06 • What Fails the Framework

A strong filter does two things: it finds good opportunities and it kills weak ones quickly. Most idea filters are optimistic by design. This one is adversarial by design. If an idea fails two or more of the six elements, it is not a pivot candidate — it is a different idea entirely.

Idea type	Elements present	What’s missing	Why it fails
Thin AI wrapper	New capability	Hard constraint, nested conflict depth	Anyone with API access can replicate it in a weekend. No moat, no depth, compressed to zero margin in 18 months.
Cool demo	New capability	Timing urgency, buyer urgency	The technology works, but no one is in enough pain right now to pay, switch, or embed it in their workflow. A solution without a customer who feels the problem acutely.
“AI for X” pitch	New capability	Old tradeoff articulated, truth-conflict identified	The pitch names a category, not an opportunity. “AI for legal” is not an idea. “The tradeoff lawyers accept between speed and accuracy, and why trust makes naive AI automation dangerous here” — that is an idea.
Broad platform from day one	Nested conflict depth (imagined)	Concrete wedge	No proof point, no first customer, no specific enough problem to be winnable. Platforms are discovered, not designed upfront.

Idea type	Elements present	What's missing	Why it fails
Regulated-market fantasy	Hard constraint (recognized)	Credible path through the constraint	The founder sees the moat but has no relationships, no regulatory experience, and no domain expertise. The constraint protects incumbents, not the new entrant, unless you can credibly operate inside it.
Timing-dependent idea, early	Old tradeoff, truth-conflict	Right timing	The opportunity is real but the capability threshold has not been crossed yet. Customers are not in enough pain yet. The window is not open.
Saturated category	All elements present	But window has closed	The opportunity structure was real — and ten well-funded companies already entered. The six elements being present does not guarantee a window is still open. Competitive dynamics are a separate check.

Insight

The framework's job is not to confirm that an idea is interesting. It is to surface which element is missing — because the missing element is usually the reason the idea has not worked for the people who tried it before.

07 • Worked Examples

Example A — AI Education (“The Formation Problem”)

Element	Analysis
Old tradeoff	Personalised learning requires a 1:1 human tutor most families cannot afford. Schools optimise for the median student.
New capability	AI can now generate adaptive explanations, practice, and feedback at near-zero marginal cost per student. The tutoring threshold has been crossed.
Truth-conflict	Children need personalised help — AND — they need discipline, shared standards, human mentorship, and moral formation. Naive AI tutors solve the first and violate the second.
Synthesis	AI handles adaptive practice and feedback. Humans handle formation, relationship, accountability, and wisdom. A deliberate division of labor, not a replacement.
Timing	AI quality crossed the tutoring threshold in 2024-25. Homeschooling is at a 20-year high. Remote learning normalised async workflows. Parents are more open to non-institutional education than at any point in recent history.
Constraint	Child safety, parental trust, pedagogy quality, worldview alignment, curriculum standards, measurable learning outcomes, regulatory compliance. These are not features — they are entry barriers.
Nested conflict	Level 1: personalized tutoring. Level 2: how does this develop discipline and deep thinking, not just correct answers? Level 3: what kind of person is this education forming — and who decides? Each level is bigger than the last.
Wedge	AI mastery tutor for homeschool families in one subject (math or reading), with parent dashboards and human-designed curriculum. One segment, one subject, one proof point.

Element	Analysis
---------	----------

Example B — AI Healthcare Administration (“The Documentation Trap”)

Element	Analysis
Old tradeoff	Doctors can either spend time with patients or keep up with documentation, billing, and compliance. The time constraint is structural — not solvable by working harder.
New capability	AI can now summarize visits, draft clinical notes, process forms, and structure patient data accurately enough to relieve the administrative burden without requiring physician replacement.
Truth-conflict	Healthcare needs automation to scale — AND — it requires human accountability, clinical judgment, and patient trust. Systems that automate too much face liability and trust collapse. Systems that automate too little don’t move the needle.
Synthesis	AI handles documentation, routing, and administrative preparation. Clinicians retain final clinical judgment and patient relationship. Clear division of labor with explicit accountability.
Timing	Clinician burnout is at a measured record high. Staffing shortages are structural. Reimbursement complexity has increased. AI documentation tools passed the accuracy threshold for clinical use in 2024. The regulatory pathway is now clearer than in 2022.
Constraint	HIPAA, malpractice risk, EHR integration requirements, clinical workflow embedding, accuracy at medical grade, procurement cycles, and physician trust. None of these are shortcuttable. Each is a compounding barrier.
Nested conflict	Level 1: documentation burden. Level 2: care coordination and handoff quality. Level 3: preventive care and chronic disease management. Level 4: outcome accountability and value-based care. Each level is structurally larger than the last.
Wedge	Specialty-specific AI admin assistant for one high-burden clinical workflow — e.g., orthopedic pre-auth and visit documentation — in a single health system. Prove workflow embedding before expanding.

08 • Hidden Principles

These principles apply regardless of which specific opportunity is being evaluated. They are generative — they produce the same insight in a situation you have never seen before.

Principle 1 — Capability Abundance Shifts Value to Judgment

When a resource becomes abundant, the bottleneck shifts to whatever is needed to use it well. Oil abundance created demand for combustion engines. AI capability abundance creates demand for judgment — the capacity to direct AI, evaluate its outputs, and decide where its application creates real value versus noise. The moat is not AI access. The moat is accumulated judgment — and the track record that proves it.

Practical implication: evaluate founders on the quality of their judgment, not the sophistication of their tech stack. The tech stack will be replicated within 18 months. The judgment compounds.

Principle 2 — Constraints Are Not the Enemy of Invention

A common mistake: treating AI opportunity as inversely related to difficulty. This produces a rush toward the easiest applications — those with no regulatory friction, no trust requirement, no workflow integration challenge. These are also the ones that compress toward zero margin fastest. Where AI capability is high and real-world constraint is high simultaneously, serious invention becomes possible and defensible. The startup that enters a constrained domain and develops the capability to operate inside it — regulatory relationships, trust infrastructure, domain expertise — builds something that cannot be replicated by an API call.

Practical implication: the best AI opportunities are in the domains that most AI founders avoid. Not because they are more virtuous, but because the constraint creates protection that easy applications cannot have.

Principle 3 — The Wedge Predicts the Platform Only If the Conflict Chain Is Deep

A wedge is a specific, winnable first problem that proves the thesis before expansion. But not all wedges predict platforms — only those where the first conflict opens into a deeper system large enough to support a company at scale. The test: after solving the first problem, does a more important problem become visible? If yes, the wedge is an entry point to a platform. If no, the wedge is the whole business — and it may be too small or too easy to protect.

Practical implication: map the conflict chain before committing to the wedge. The wedge is the right starting point only if the chain beneath it is deep enough to matter.

Insight

The three principles converge on a single insight: in the AI age, the question “can we build it?” has been answered affirmatively for almost everything. The frontier question is “does the structure beneath this opportunity — the tradeoff, the conflict, the timing, the constraint, the depth — make it worth building?” That structural question is not answered by a better AI model. It is answered by better judgment.

Case Study — Framework Applied: PASS / WATCH / KILL

Full framework applied to one concrete startup category. Company is hypothetical; category is real.

The Idea: AI Prior Authorization and Clinical Documentation for Specialty Care

One-line description: An AI system that handles prior authorization requests and clinical note documentation for specialist physicians — reducing the administrative time that forces the tradeoff between patient volume and care depth.

Element	Analysis	Strength
---------	----------	----------

Six-Part Formula Applied

Element	Analysis	Strength
Old tradeoff	Specialty physicians (orthopedics, cardiology, oncology) spend 15–20 hours per week on prior auth and documentation — time structurally taken from patient care. The tradeoff is not solvable by working harder; it is embedded in the reimbursement system.	Strong
New capability	AI can now draft clinical notes, structure visit summaries, pre-populate prior auth forms, and generate appeal letters at accuracy levels acceptable for physician review — not for autonomous submission, but for physician-supervised completion. The accuracy threshold was crossed for structured clinical documentation in 2024.	Strong
Truth-conflict	Healthcare needs workflow automation to scale and reduce burnout — AND it requires physician accountability, patient trust, and malpractice liability to remain with a licensed human. Naive automation (AI submitting on its own) violates the second. Physician-supervised AI preserves both.	Strong
Synthesis	AI handles drafting, pre-population, and routing. Physician reviews, edits, and approves. Accountability stays with the license holder. Speed and quality improve without removing human judgment from the liability chain.	Strong

Element	Analysis	Strength
Timing	Clinician burnout at record levels (55%+ in AMA surveys, 2024). Staffing shortages structural. Prior auth volume increased 21% in 3 years (AMA, 2024). AI documentation tools passed clinical-grade accuracy for structured workflows in 2024. CMS proposed prior auth reform creating further urgency. The regulatory pathway is cleaner than in 2022.	Strong
Hard constraint	HIPAA compliance, BAA requirements, EHR integration (Epic, Cerner, Athena), malpractice risk architecture, payer-specific prior auth formats (each payer has different portals and criteria), physician trust and change management, clinical accuracy requirements. None are shortcuttable. All compound as moats for the operator who navigates them.	Very strong
Nested conflict depth	Level 1: documentation burden. Level 2: prior auth denial rate and appeal workflow. Level 3: care coordination between specialists and primary care. Level 4: chronic disease management and remote monitoring. Level 5: value-based care reporting and outcome accountability. Each level is structurally larger than the last.	Very strong

Buyer and Commercial Path

Item	Analysis
Primary buyer	Specialty practice administrator / physician group CFO
Budget source	Administrative staff cost reduction; revenue cycle improvement from faster prior auth approval; physician time recovered (at \$200–400/hr loaded cost, 5hrs/week recovered = \$50–80K/yr per physician)
First paid wedge	Orthopedic prior auth + post-visit note for one specialty group or ASC. Prove: time saved, denial rate reduced, physician satisfaction.
Path to \$10M ARR	50–80 specialty groups at \$120–200K/yr each (4–6 physicians per group, \$25–35K per physician). Achievable with a 10-person team in 18–24 months post-product-market fit.

Item	Analysis
Why a startup wins	Epic and Cerner have prior auth modules that are incomplete and not AI-native. Athena is cloud-based but not specialty-optimized. No incumbent owns the AI-native, specialty-specific, physician-supervised workflow. The wedge is a gap in existing vendor coverage.

Failure Pattern Check

Pattern	Present?	Notes
Thin AI wrapper	No	Hard constraint (HIPAA, EHR, payer-specific) prevents easy replication
Cool demo	No	The pain is acute and measurable — 15-20 hrs/week lost is a known, budgeted problem
Broad platform from day one	Risk	Must resist expanding beyond one specialty until the first wedge is proven
Regulated-market fantasy	Low risk	Constraint is real, but there is a clear path through it: BAA, HIPAA-compliant infrastructure, EHR integration, physician-supervised model. Not trivial — but traversable.
Timing-dependent too early	No	All signals say now: burnout, volume, cost pressure, AI accuracy threshold crossed
Saturated	Watch	Competitors exist (Nuance/DAX, Ambience, Nabla). Market is early but moving fast — the window is 18–24 months before incumbents strengthen their AI position.

Kill Criteria — downgrade to Watch or abandon if any of the following occur:

Trigger	Threshold	Why it matters
Physician edit time exceeds manual workflow	By week 6 of pilot	If AI creates more work than it saves, the value proposition is broken before trust is established
No paying contract by day 90	At least \$2K MRR from one group	Validates willingness to pay — desire without payment is not product-market fit
EHR integration path exceeds 12 months or requires enterprise-only procurement	Any integration path longer than 12 months	Locks the product to a slow-moving market; makes unit economics unworkable

Trigger	Threshold	Why it matters
Denial-rate improvement unmeasurable at week 8	No reduction vs baseline	Prior auth ROI is the commercial anchor — if unmeasurable, the buyer cannot justify renewal
A well-funded competitor locks the first specialty wedge before launch	First-mover secured in orthopedic or cardiology prior auth	The wedge is narrow enough that second-mover disadvantage is structural in the first 18 months

Insight

Verdict: PASS — *with the condition that the first 90-day validation is completed before raising a priced round. The opportunity structure is strong on all six elements. The buyer exists, the budget is clear, and the constraint is a moat not a wall. The risk is execution: physician adoption speed, EHR integration timeline, and competitive window. None of these are reasons to kill — they are the diligence questions that the first 90 days must answer.*

SYNTHESIS

Anti-Thesis — The Strongest Case Against This Framework

The over-intellectualization problem. The six-part formula may be too deterministic. Many great startups begin with customer obsession, distribution advantage, or founder-market pull — long before the full structural picture is visible. Stripe was not built because Patrick Collison ran a six-part analysis. Airbnb was not validated by a truth-conflict table. The pattern is visible in hindsight. It is rarely the cause of the decision in the moment. A founder who waits for all six elements to crystallize before starting may wait forever.

The distribution objection. Some markets are won by speed, brand, and network effects — not by deep truth-conflict synthesis. WhatsApp had no hard constraint moat. Instagram had no nested conflict chain visible at launch. They won on distribution, timing, and product simplicity. The formula would not have flagged them as obvious.

The execution reality. A structurally elegant idea with weak execution loses to a structurally mediocre idea with exceptional execution, every time. The framework evaluates opportunity structure. It says nothing about the team's ability to operate inside the constraint, find the first 100 customers, or survive the first 18 months of nothing working.

The luck variable. Timing, in particular, contains irreducible uncertainty. A formula can identify that the window is theoretically open. It cannot tell you whether you are six months early or six months late — and that difference is often the company.

The answer to the anti-thesis:

All of this is correct. And it does not invalidate the framework. The six-part formula is a pre-execution filter, not an execution guarantee. Its job is to identify whether a specific opportunity deserves scarce founder time and capital before the team spends years proving a weak structure. It does not replace customer discovery, execution quality, or timing judgment. It narrows the field so that when those things are applied, they are applied to the right problem.

The distribution objection is the strongest. The honest answer is that some large companies do not fit the formula — and the formula does not claim otherwise. What it claims is that the companies which fit the formula, when the formula is satisfied, tend to be more defensible, more expandable, and more durable than those that do not. That is a probabilistic argument, not a deterministic one.

Insight

A framework that cannot argue against itself is not a framework. It is a belief system. The six-part formula is strongest when used as a filter that kills weak ideas quickly — not as a guarantee that passing ideas will succeed.

09 • The Synthesis

What Changes By Audience

AI compresses execution. It does not compress judgment. The gap between “this can be built” and “this is worth building” has never been wider.

For founders: stop filtering ideas by buildability. Filter by structural completeness — all six elements, each one articulable precisely. An idea that fails two elements is not a pivot candidate; it is a different opportunity entirely.

For investors: the mispriced category is constrained domains — healthcare, education, legal, regulated industries. Easy AI is already crowded. Hard AI, where the constraint is the moat, is underinvested.

For operators and executives: the nested conflict question is your strategic roadmap. After the problem your company solves today — what more important problem becomes visible? That answer is your acquisition thesis, your product roadmap, and your category-creation opportunity simultaneously.

The Remaining Truth

There is an implicit ninth question the framework does not state but always requires: *who is the buyer, and what do they already believe?* The truth-conflict framing matters only if the target customer already feels the tension acutely enough to act. If they have accepted the tradeoff without distress, the opportunity may be real but the timing is early. The framework is a filter for opportunity structure. The market is a filter for the moment. Both must be present simultaneously.

Closing

The average AI founder will build what is newly easy. The great AI founder will build what is newly possible but still deeply hard.

That is not a moral distinction. It is a strategic one. What is newly easy will be crowded, cheap, and commoditized within 24 months. What is newly possible but still deeply hard — because the tradeoff is real, the conflict is genuine, the timing is precise, the constraint is high, and the problem chain is deep — is where durable companies are built.

The six-part formula is a tool for finding the latter before the market recognizes it. In the AI age, that window is narrow and accelerating. The judgment to see it clearly is the most valuable capability a founder or investor can have.

CHECKLIST — FINDING THE REAL ONES

Insight

How to use in investment diligence: Apply the six-part formula before market sizing. If the old tradeoff, timing, constraint, and wedge are vague, do not proceed to TAM, valuation, or comparable analysis. The opportunity structure is not yet understood — and no amount of market-sizing will fix a weak structural foundation. The formula is a pre-diligence gate, not a diligence step.

Apply this to any startup idea. One page. Adversarial by design.

Part A — The Six Elements (structural test)

Element	Your answer	Strength: Weak / Moderate / Strong
Old tradeoff — what painful compromise has the market accepted as permanent?		
New capability — what specific recent change makes this tradeoff breakable now (not “AI is improving”)?		
Truth-conflict — what two independently true things are in tension? Does naive AI violate one of them?		
Timing — why now and not two years ago? Name the specific change that opened the window.		
Hard constraint — what makes this defensible against a well-resourced competitor in 12 months?		
Nested conflict depth — after the first problem is solved, what more important problem becomes visible?		

Structural verdict: If two or more elements are Weak → the opportunity structure is not yet real. Do not proceed until they strengthen.

Part B — The Failure Pattern Check

Failure type	Does this apply?
Thin AI wrapper — no constraint, no depth	Yes / No
Cool demo — no buyer urgency right now	Yes / No
“AI for X” pitch — no tradeoff or truth-conflict articulated	Yes / No
Broad platform from day one — no concrete wedge	Yes / No
Regulated-market fantasy — constraint present but no credible path through it	Yes / No
Timing-dependent idea, early — window not yet open	Yes / No
Saturated category — structure was real but window has closed	Yes / No

If any apply: name which one, and name the specific fix required before proceeding.

Part C – The Buyer Test (the implicit ninth question)

Does the target customer already feel the truth-conflict acutely — or have they accepted the tradeoff without distress?

If they feel it acutely: timing and market urgency are aligned. Proceed.

If they have accepted the tradeoff: the opportunity may be real but the timing is early. Do not confuse a future opportunity with a current one.

Sources

[1] Epoch AI, *Tracking AI Cost Trends*, cost-per-token series (epochai.org/data/trends); OpenAI API pricing changelog, January 2023 – January 2025.

[2] a16z, *State of AI 2024* (a16z.com/ai); Crunchbase, *AI Startup Funding Trends Q1–Q4 2024*.

[3] Microsoft, *Work Trend Index 2026: AI at Work* (microsoft.com/en-us/worklab/work-trend-index); McKinsey & Company, *The State of AI: Global Survey*, April 2024.

[4] European Commission, *EU Artificial Intelligence Act — Official Timeline* (ai-act-service-desk.ec.europa.eu); EUR-Lex, Regulation (EU) 2024/1689.

[5] Goldman Sachs Global Investment Research, *AI Investment in the Physical Economy*, June 2026; Axios, *AI Capital Tracker*, H1 2026.

[6] Stanford HAI, *AI Index Report 2026*, Chapter 4: Safety and Alignment (hai.stanford.edu/ai-index); *International AI Safety Report 2026* (internationalaisafetyreport.org).

[7] Tseng, P. et al., *Administrative Costs Associated with Physician Billing and Insurance-Related Activities*, JAMA Health Forum, 2022; MGMA, *Cost Survey for Physician Practices*, 2024.

Disclaimer

Not investment advice. This report is a structured thinking study produced by Lantern Labs as part of its investment research practice. Nothing in this report constitutes investment advice, a recommendation to buy or sell any security, or an offer or solicitation to invest in any fund, vehicle, or company. It should not be relied upon as the basis for any investment decision.

Illustrative cases only. All company examples, worked examples, and scenario analyses in this report are illustrative. They are used to demonstrate the application of a conceptual framework, not to make claims about specific named organisations. References to companies such as Ford, Apple, Amazon, Airbnb, and Stripe are used for educational illustration only and do not constitute research coverage, endorsement, or analysis of those companies' securities or prospects.

Forward-looking statements. The 2036 Hindsight Test and related sections contain forward-looking statements and hypothetical scenarios. These reflect structured analytical exercises, not forecasts, predictions, or investment theses. Actual outcomes will differ materially from any scenario described here.

No warranty. Lantern Labs makes no representation or warranty, express or implied, as to the accuracy, completeness, or fitness for purpose of the information contained in this report. The framework presented is a decision aid, not a guarantee of outcome. The six-part formula is a pre-execution filter. It does not predict which companies will succeed.

Confidentiality. This report is prepared by Lantern Labs for its clients and contacts. It is not intended for public distribution unless explicitly authorised. Recipients should not reproduce, redistribute, or share this report without prior written consent from Lantern Labs.

Lantern Labs | www.lanternlabs.net | Decision Intelligence, not Decision Making

July 2026 · Not investment advice · Illustrative cases — not named organisations