AEIOU ACADEMY

STRUCTURAL LITERACY SERIES

Making the Go/No-Go Decision

The Four Gates Test

Four questions, in fixed order, that can be answered before capital is committed

An analysis for founders, boards, and investors building in regulated, high-consequence technology. Surgical robotics is the laboratory here, not the subject.

Christopher J. Velis, NACD.DC

Founder, Auris Health · Co-Founder, AEIOU Academy

Critical review: Dr. Christos Kelepouris, Co-Founder, AEIOU Academy

July 2026 · WHITE PAPER

Series Note

Contributions: CJV conceived the framework, conducted the analysis, and wrote the paper. Dr. Christos Kelepouris provided substantive critical review across successive drafts; his challenges materially changed several claims and the sourcing behind them. Responsibility for the argument and for any remaining error is the author's alone. Not legal, financial, or investment advice.

This paper belongs to the Structural Literacy Series published by AEIOU Academy. Part 1, *Capturing Value at Exit*, and Part 2, *The Structural Alternative*, examine what happens to founder value at the moment of sale. This paper works the other end of the same problem: whether the venture was ever worth building, and how that question can be answered before capital is committed rather than reconstructed afterward.

Each paper is self-contained.

How to Use This Paper

Founders and CEOs. Start with Part I and Part II, then run the diagnostic in Part XII against your own company. The test takes about an hour and requires nobody's permission.

Boards and investors. Part XII sets out the staging rule: a failed first gate is a decline, an unanswered one is a smaller investment in the answer itself. Part VI applies the framework across seven companies; Part IX states what it cannot reach, including a structural mismatch between medical devices and public equity that affects how these companies should be financed.

Founders who have already failed the first gate. Part XII opens with the four moves available and the distinction between a company that never passed and one that passed at founding and would fail today. The remedies differ.

Anyone outside medtech. The section immediately below explains why this field produces unusually clean evidence about failures that occur everywhere.

Why read this

Vicarious Surgical was well funded and hired well. Its patent portfolio was strong enough that Intuitive Surgical, which has held the dominant position in this industry since the da Vinci system cleared in 2000, cited it forty-four times while designing around it.

In July 2026 the company dissolved, having consumed $253.4 million since inception. It had never sold anything.

Others went the same way. Medrobotics filed for Chapter 7 liquidation with its cleared system already installed in thirty hospitals; Hansen Medical was acquired for $4.00 per share, roughly $80 million; Titan Medical failed to sell the business after approaching more than fifty parties, licensed its robotics intellectual property to Intuitive, and merged into a cardiovascular imaging company; Asensus accumulated $987.6 million in losses and was acquired for roughly a tenth of that.

Ask why and you get the same list every time. The technology was harder than expected. The burn rate was unsustainable. The capital markets turned. Competitors moved faster. Each is vague at best,

and none gives much insight into what actually happened. Every item on that list also describes Intuitive Surgical between 1995 and 2000, and Intuitive went on to own the field. A diagnosis that cannot separate the living from the dying is not a diagnosis.

This paper proposes four conditions that determine whether a capital-intensive technical venture survives, and argues they resolve in one direction only. The first condition governs the rest. The more useful claim is the second one: each condition can be tested against public records before the money is committed, on dates anyone can look up.

None of these failures came from bad engineering. Most of the machines worked. The people involved were experienced, well advised, and well funded, and they answered the questions out of order.

If you build, fund, or sit on the board of a device company, the test in Part II takes about an hour and needs nobody's permission.

Why surgical robotics is the right laboratory

Anyone working in climate technology, autonomous systems, aerospace, diagnostics, advanced manufacturing, energy, or defense will be tempted to skip a paper about surgical robots. The reason not to is methodological.

Four properties rarely occur together: an external regulatory gate, an absolute reliability requirement, frontier talent and capital, and a market still forming. Surgical robotics has all four, which makes it an unusually clean place to see what determines outcomes.

The field is regulated, so the gate to market sits outside the company, on a public record, with dates attached. You ship when the agency says so.

The product cannot fail, because the failure mode is a patient on a table. Engineering gets held to a standard most industries never reach, so when one of these companies says the machine worked, the bar was high and it was usually cleared.

The work sits at the technological frontier and attracts first-rate engineers and serious money. Nobody in this record ran out of talent.

And the market is still forming rather than divided, so the question of what a customer would otherwise do stays open instead of answering itself.

Put those together and the usual explanations stop being available. A venture that fails here was not underfunded, was not badly engineered, was not run by amateurs, and was not scrapping for share in a mature market. Strip those away and what remains sits close to the actual cause.

The field therefore repays attention from people who will never see an operating room. It leaves a cleaner record than almost anywhere else, and the questions it settles are the ones you face.

The framework transfers directly to capital-intensive ventures with long development cycles, institutional buyers, complicated adoption, and a need to capture recurring value. Elsewhere it needs adaptation, and Part XII says where.

Abstract

Capital-intensive technical ventures fail at a high rate, and the explanations offered afterward are usually indistinguishable from the conditions the survivors also endured. Engineering difficulty, capital markets, and competitive intensity describe winners and losers equally well, which means they explain nothing.

This paper argues that outcomes are governed by four conditions in fixed order. **Substitutes:** is there an outcome no other path achieves, and can a machine deliver it. **Right to practice:** can the company commercialize without the incumbent's permission, and keep an economically meaningful share of what it creates. **Demand:** how many units, at what price, with what feature set. **Capture:** does each installed system produce recurring revenue.

The order carries the argument. The possibility of value-based pricing is established at the first gate, units at the third, and the recurring multiple at the fourth, while the second determines whether any of it is yours to keep. A rich consumable stream multiplied by one placed system is nothing, and a patent portfolio protecting a capability nobody needs protects nothing.

The central claim is not retrospective. These conditions are assessable in advance, and the industry's standard practice is not to assess them but to substitute proxies for them: patent counts, team pedigree, and category-level market size. Vicarious Surgical's 2026 dissolution is examined against its filing record to show that the determining question was answerable at incorporation in 2014 and appears not to have governed a financing decision in the twelve years that followed. Two counter-cases bound the framework: a company that cleared the first gate and went bankrupt anyway, and a system that fails the first two and thrives because its parent already owns the annuity.

Surgical robotics supplies the evidence because it is regulated, held to a near-zero defect standard, technologically advanced, and commercially emergent, a combination that excludes most of the ordinary excuses and leaves the causes unusually exposed. The findings apply to capital-intensive technical ventures generally, with the boundary conditions stated explicitly.

The paper closes on what the gates cannot reach, including a structural mismatch between medical devices and public equity, and on a correction to how founding risk is usually described. A founder does not inherit the irreducible risk of a venture. He selects it.

How the paper is structured:

Part I dismantles the standard post-mortem and answers the hindsight objection. Parts II through V build the four gates and explain why the order carries the argument. Part VI applies them across seven companies, Part VII to Vicarious against its filings, and Part VIII to two cases that falsify simpler versions of the thesis, one of them outside surgery entirely. Part IX states what the gates cannot reach. Part X answers the remaining objection, that the outcomes were luck. Parts XI and XII place the framework against Porter, Christensen, and Moore, and say what to do if you fail the test.

Part I: The List That Explains Nothing

On 21 July 2026, the shareholders of Vicarious Surgical voted to dissolve the company and hand the assets to a liquidator.

Twelve years. Two hundred forty million dollars. Twenty-six people left at the end. No revenue in any fiscal year since incorporation.

The next morning the FDA authorized Johnson & Johnson's OTTAVA surgical robot for ten general surgery procedures, with a hernia study already running. One company went to the liquidator and the other went to market, in the same anatomy, in the same week.

So the category was alive. Something else had died.

Everyone already knows why it failed

Ask around and the answers arrive in roughly this order. The technology proved harder than anyone expected. It never reached commercialization. The burn rate ran on against no revenue. The company went public through a special purpose acquisition company, or SPAC, a listed shell that takes a private business public by merging with it, and then the capital markets turned. Competitors kept advancing while it was still building.

All five are true, which is the problem.

Apply them to a different company. Founded in 1995. Five years with no cleared product and no revenue while a rival with an already-cleared robot took the early market. Venture money burning the whole time. Unprofitable for most of a decade, through the dot-com collapse.

That company is Intuitive Surgical, which defined the field and reported $10.1 billion of revenue in fiscal 2025.

Every item on the list applies to both. The list describes what building a surgical robot feels like, which is useful if you want sympathy and useless if you want a decision.

The question the industry asks

Walk a medtech conference floor and you hear one question in fifty variations. How do we build a better surgical robot?

More degrees of freedom. Smaller ports. Better ergonomics, haptics, visualization. Lower cost. Now, inevitably, software and AI. Every entrant arrives with a superior answer and a deck explaining why the incumbent's architecture is a decade old.

Good engineering question. Wrong commercial one.

The commercial question is what operation becomes possible. Almost everything follows from that answer, and almost nothing follows from the other.

One definition before we start

A surgical robot is not a market.

The unit of analysis in this paper is the anatomy-indication pair: a specific procedure, in a specific anatomical region, against its own specific substitutes. Peripheral lung nodule biopsy is a market. Ventral hernia repair is a market. "Surgical robotics" is a category description, and analysis at that level produces answers true of nothing in particular.

Outside medicine the equivalent is the use-case-constraint pair: the specific job, in the specific setting, against the specific alternative the customer would otherwise choose.

The distinction carries more weight than it appears to. Foreclosure runs per anatomy rather than across a field. A multi-specialty platform has to pass the test separately for every indication it claims. Two companies building similar machines can face opposite verdicts.

Nearly every error in this paper begins with analysis performed at least one level too high.

Four questions, and they run in order

Figure 1. The four gates. The equation says any zero is fatal; the order says it is not optional.

One. Substitutes. Is there an outcome no other path achieves, and can a machine deliver it?

Two. Right to practice. Can you commercialize without the incumbent's permission, and keep an economically meaningful share of what you create?

Three. Demand. How many units, at what price, with what feature set?

Four. Capture. Does each installed system produce recurring revenue?

Value-based pricing becomes possible at the first gate. Units are decided at the third, the recurring multiple at the fourth. The second determines whether any of it stays yours.

Fail the first and the rest are decoration. A rich consumable stream multiplied by one placed system comes to nothing. A patent portfolio protecting a capability nobody needs protects nothing.

What "fails" means here

The claim depends on the definition, so here it is.

Failure in this paper means no financial return, or a return too small to justify the capital, together with no durable independent survival. Success means an independent company at venture scale whose economics rest on its own product.

Gate 1 is dispositive for that outcome and no other. A technology that fails it can still be sold to an incumbent as a defensive instrument, and Part IX examines a system that fails the first two gates and does well inside a parent that already owns the fourth. Both are real outcomes. Neither is the one this paper is about, and every use of the word dispositive should be read against that boundary.

Why this is not hindsight

The interesting claim is not that Vicarious failed. It is that the answer sat in public, on a date you can look up.

Vicarious was building a single-port robot for abdominal soft tissue, with ventral hernia repair as the announced beachhead. A word on what clearing that market requires: in the United States most devices reach it by showing the FDA that they are substantially equivalent to something already sold, a route called a 510(k), and the device they are compared against is the predicate. Part III returns to this, because the choice of route turns out to say a great deal about a company's position. Surgeons perform ventral hernia repair laparoscopically thousands of times a day. They were doing so in 2014, when the company incorporated, and Intuitive's multiport da Vinci had been doing it robotically for years.

No moment existed when the substitute was absent. Not at founding, not at the shell merger, not at the end. Checking required no forecast and no inside knowledge.

The company raised capital for another eight years, including a reverse merger with that shell, which valued it near $1.19 billion in September 2021, the same month a private round priced it at $98 million. Twelvefold in thirty days, with nothing clinical or regulatory changed. The venue changed.

And the part worth sitting with: Vicarious held 39 patent families and 63 active grants, and Intuitive cited that portfolio forty-four times. The market leader was reading those patents and designing around them.

They had the intellectual property. It could not save them, because the gate it satisfied sits downstream of the gate that failed.

Part II: The First Gate

Start with the lung

A nodule sits in the far periphery of the lung, down an airway branch too narrow and too crooked for a bronchoscope, the then standard of care, to follow.

You still have ways to reach it. Open the chest. Put a needle through the chest wall. Push a catheter blindly past the end of the scope and hope it lands somewhere useful.

All three get you tissue. None gets you tissue without opening the chest, puncturing the pleura, or working blind, and the whole of this paper hangs on that difference.

The test is about outcomes, not procedures

The gate asks nothing about competitors, and nothing about whether the job can be done another way. Almost everything can be done another way. It asks whether the **outcome** can be obtained another way, at the same cost to the patient.

A thoracotomy is an alternative, not a substitute, because it charges the patient a chest incision the other path does not.

Specify the outcome loosely and the substitute set runs to infinity. Specify it tightly and it collapses to a handful, or to none. Founders specify loosely, because the loose specification is the one under which their product looks necessary.

Two questions, both yes or no

1A. Is there an outcome no substitute achieves?

1B. Can a machine deliver it?

No units, no price, no feature analysis. Those belong to the demand equation at Gate 3, and running them before this gate is settled means working on a company that has not earned the attention.

What "the same outcome" means

Left undefined, "outcome" becomes whatever the founder needs it to be. It runs on four dimensions.

Clinical result. Does the alternative reach the same medical end? The diagnosis obtained, the tumour removed with clean margins, the joint replaced and working.

Patient burden. At what cost? An incision, blood loss, nights in hospital, weeks of recovery, permanent loss of function.

Access. Can this patient get it at all? Geography, surgeon availability, and the population currently receiving nothing.

Reproducibility. Can a trained practitioner deliver it reliably, or only an exceptional one?

The gate turns on whether the best available alternative fails **decisively** on at least one dimension. Decisive means a step change a clinician would describe categorically, not a percentage. The standard everywhere is "cannot," never "better."

Companies that fail here usually improve all four a little and none of them much. What they have built is a product. A hospital will not replace working equipment and retrain trained surgeons for gains on every axis and a transformation on none.

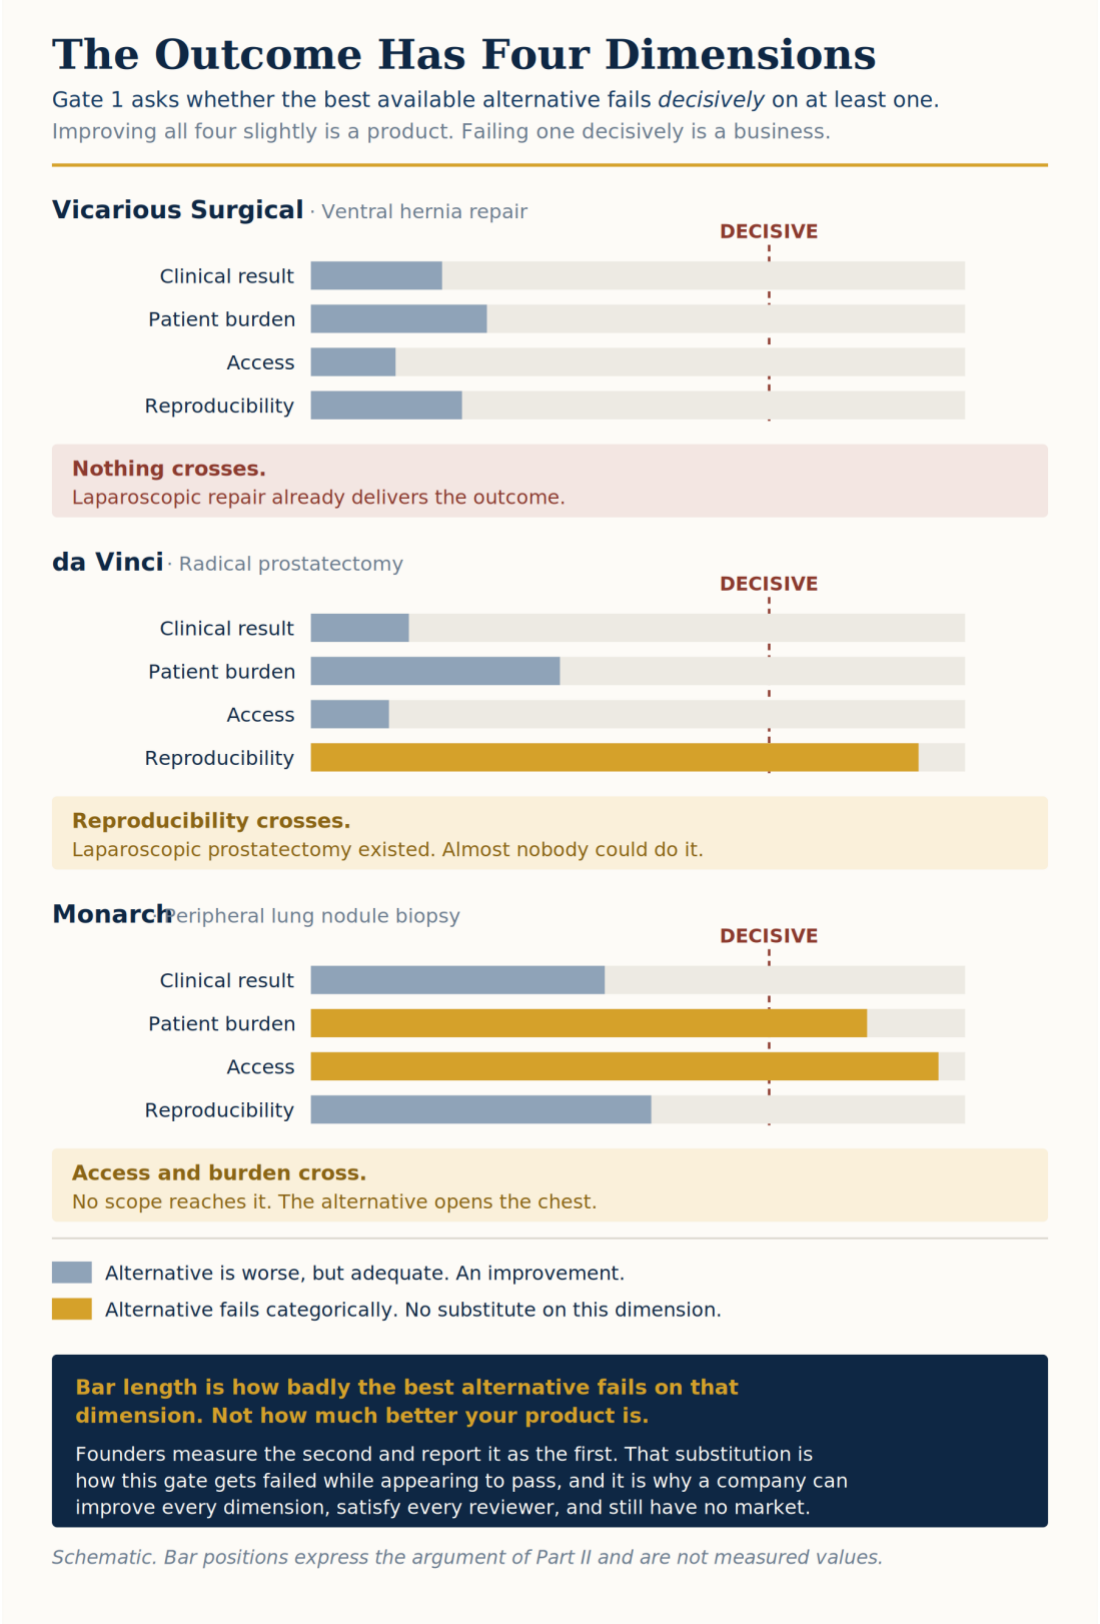

Figure 2. The outcome has four dimensions. Gate 1 is a threshold test, not a comparison.

The dimension the industry under-reads

Reproducibility gets missed, and missing it obscures how the most valuable company in the field was built.

Laparoscopic radical prostatectomy existed before da Vinci. The clinical result was comparable and the patient burden ran far below open surgery. Hardly anyone could perform it. The operation demanded exceptional laparoscopic skill in a confined pelvic field at awkward angles, and the surgeons who could do it reliably numbered in the dozens.

Da Vinci did not make the operation possible. It made the operation reproducible, by ordinary surgeons, at population scale, which is a decisive failure of the alternative on one named dimension and a clean pass at Gate 1.

Now the argument people usually make instead.

Robotic prostatectomy scaled to 85 percent of the American procedure volume before any randomised trial had compared it to open surgery, and when the trial reported it found no difference in urinary or sexual function. Patients went to the robot center regardless and treated it as the center of excellence. The perception was real and it moved adoption fast.

Perception is a demand strategy, and demand is Gate 3.

Holding those apart matters more than it looks. A founder who believes perception clears the first gate will build a marketing program for a capability that has substitutes, and spend years learning that the hospital was never the party needing persuasion. Intuitive had a real answer on reproducibility and ran a perception strategy on top of it. Two decisions, and only the first one was Gate 1.

The access dimension, and the trap inside it

The gate can also be cleared without widening the procedural envelope at all.

Telesurgery invents no operation. It puts an existing operation within reach of a patient who could not otherwise have it, because no surgeon is within range. For that patient no substitute exists, and no hand reaches across an ocean.

SS Innovations has performed cardiac telesurgery across roughly 20,000 kilometres of network distance at latency near 300 milliseconds, which answers the first gate.

What it leaves unanswered is who pays, and the problem surfaces hard at Gate 3.

What passing 1A buys

The possibility of value-based pricing. Not the price.

Where nothing else reaches the outcome, price can be set by value delivered. Where close substitutes exist, price is set by rivalry, and rivalry in capital equipment is a fight the newcomer loses on arithmetic: unamortised development, fewer systems already placed in hospitals, no service scale, no consumable volume. The incumbent prices to defend. The challenger prices to survive.

Which explains why a low price gets misread so often. Presented as strategy, it is usually a symptom, and what it is a symptom of is substitutes.

How much of that potential survives reimbursement, capital budgets, procurement and contracting belongs to Gate 3, and Gate 3 takes a great deal of it away.

Your FDA filing is evidence, not the answer

The regulatory pathway a company chooses tells you something, provided you do not ask it to carry more than it can.

A 510(k) requires demonstrating substantial equivalence to a legally marketed predicate: same intended use, and either the same technological characteristics or differences raising no new questions of safety and effectiveness. A de novo request asserts that no legally marketed device offers a suitable basis for the comparison.

Be careful about what those findings mean. Substantial equivalence describes a regulatory relationship between two devices. It says nothing about whether the predicate delivers the same clinical outcome, the same patient burden, or the same economics, and the predicate need not still be on the market. De novo likewise establishes nothing about whether a procedural, pharmacological, or manual alternative exists. The agency answers its own question, not yours.

What the pathway does is impose an obligation to explain.

A 510(k) means you have told a federal agency that something already marketed is substantially equivalent to your device. You now owe your board an account of why regulatory equivalence is not commercial substitutability.

Sometimes a good answer exists. Often the attempt to write one is where a founder discovers it does not. SS Innovations chose a 510(k) over a de novo for the SSi Mantra in December 2025, after a pre-submission meeting, and took on that obligation. Johnson & Johnson, taking OTTAVA through de novo, did not.

What I chose, and why

A word about where I am standing.

I incorporated Auris Surgical Robotics in Delaware in 2007 and ran it as sole director and chief executive until 2011, when I recruited an operator to scale it. The company was built to reach anatomy no hand could reach, and choosing to aim there, rather than at a better version of something surgeons already performed, was the first real decision it made. There was no machine yet to argue about.

I did not want a contested-superiority fight, or the comparative studies needed to win one. I wanted a capability so plainly beyond the human hand that nobody would think to argue about it.

The choice concerns the shape of the critical path more than the shape of the product. Put your first gate inside an evidence argument and you have handed the schedule to a process you do not control, with your opponent selecting the endpoints.

Nobody can benchmark you

Having no substitute carries a second benefit, separate from the commercial one, and it goes largely undiscussed.

If nothing comparable exists you cannot be compared. Diligence has nothing to measure you against, so you get assessed on your own claim rather than against another company's numbers. The advantage shows up in financing and shows up larger at exit, because a strategic buying a differentiated asset is not running a bake-off.

Absence of a comparator and absence of a competitor are different things. The first is an information condition, the second a market condition. Have the second without the first and diligence arrives with a spreadsheet.

1B: can a machine do it

Engineering risk and biological risk sit together here, because building is the only way to resolve either and neither yields to research. They behave differently and they fail differently.

Engineering risk yields to iteration and, more to the point, to capital. Enough cycles and enough talent and mechanisms converge. Expensive, slow, tractable. We were not guessing how biology would respond. We were building a machine, like any other machine.

Biological risk yields to neither with any reliability. Tissue, anatomy and physiology are not negotiable design constraints. A mechanism that performs on the bench can be defeated by the material it operates on, and the failure surfaces late, after the money is spent and the team is built.

1B is the only irreducible element in this framework.

Substitutes, demand, right to practice and capture all yield to research, counsel or arithmetic before a dollar of engineering capital is committed. The build alone has to be attempted before it can be known, and Part X is built on that fact.

Do them in order

Settle 1A. Then take 1B.

In that order, engineering risk gets taken against a question already answered. Invert it and you take engineering risk first, then go looking for the indication, which reliably produces an impressive platform aimed at a procedure already served well enough.

Vicarious filed its first patent on 5 May 2014, the year it incorporated, under the title "Virtual reality surgical device." The founding artifact was an architecture. The indication arrived years later, chosen to fit a machine that already existed.

The inversion is not a process flaw. It answers the first gate before the company notices the question has been asked.

The gate moves while you build

Here the trap catches good companies.

You do not run the substitute test against today's field. You run it against the field on the day you can ship, and the regulatory pathway sets how far ahead you must see. A long pathway means your answer has to hold for five years or more, and often it does not.

The trap is not particular to robotics. It ran to completion in motion preservation of the spine after Charité and ProDisc, where billions went in and the category had already closed. Patents were not the barrier. Years were. By the time a follower could ship, someone else had settled the question.

Substitute status decays after clearance too. A capability with no substitute at founding acquires one. The decay is why the other three gates exist, since they are what you hold once the first stops being exclusive.

Part III: The Second Gate

A letter arrives

You have done everything right. The capability is real, the machine works, and the FDA has cleared it.

Then a letter arrives from the market leader, listing the patents you infringe.

Nothing stops. A patent confers a right to exclude, and the right gets exercised through litigation that runs for years, so you keep shipping until a court enjoins you, until you settle, or until the cost of defending yourself finishes the company.

The last one is the mechanism that matters, and it has nothing to do with the merits. Your opponent is spending a rounding error. You are spending the next round. Whoever holds the deeper balance sheet is usually holding the pen when the settlement gets drafted, and the settlement is where your margin and your freedom of movement go.

Can you commercialize without the incumbent's permission?

Two questions wear that one label. **Freedom to operate** asks whether you can legally sell. **Appropriability** asks whether you can keep an economically meaningful share of what you create. A company can pass one and fail the other.

Three ways through

License it. Fastest, and it caps you. Your margin and part of your strategic freedom now belong to somebody who would prefer you did not exist.

Build a portfolio that makes the fight unattractive. Not a portfolio that wins. A portfolio expensive enough to attack that the incumbent would rather leave you alone or buy you. The classic route, and the one that makes a company acquirable by a strategic wanting into the category without inheriting a lawsuit.

Take a regulatory position the incumbent does not hold. New, and taken up below.

What that looked like in practice

Computer Motion had the first FDA-cleared robot for abdominal surgery. It did not lose in the operating room.

Computer Motion sued first, in the Central District of California in May 2000, over nine of its own patents. In April 2001, IBM and Intuitive Surgical sued jointly in Delaware over an IBM patent Intuitive held under exclusive license. A Delaware jury found Computer Motion liable and awarded $4.4 million. The California case was still awaiting trial in 2003.

Both companies bled. In 2003 they merged, and the pioneer ended up inside its competitor.

Now look at what Intuitive had already done. Its December 1997 license from IBM ran exclusive in the field and covered two groups of patents: the LARS patents, relating to augmentation of endoscopic surgery, and the ROBODOC patents. The agreement gave Intuitive the right, though not the obligation, to enforce them at its own expense.

Three years before da Vinci cleared, Intuitive had bought not a position but the right to enforce a foundational estate it had no hand in creating. Almost nobody takes this gate that early.

The gate has been quietly commoditized

Here is a fact that explains the current wave of entrants better than any story about breakthrough engineering.

The foundational mechanism patents have aged out. Twenty years from filing, and the filings that built the original platform date to the 1990s and early 2000s. Medtronic, CMR, Distalmotion, SS Innovations and others could all arrive within a few years of each other, not because the technology suddenly became possible but because the wall came down on schedule.

Time did what none of them could do.

The new route: build the predicate yourself

A de novo authorization does more than clear a device. It creates a classification.

When Johnson & Johnson took OTTAVA through de novo in July 2026 as the first table-integrated soft tissue robotic system, it did not only earn permission to sell. It became the predicate, and anyone following into that architecture now demonstrates substantial equivalence to J&J's device.

The two gates connect here in a way worth noticing. A real answer at the first gate is what makes a predicate impossible to find, and no predicate is what qualifies you for de novo. **Clearing the first gate makes the second one cheaper.**

Where artificial intelligence acts

The consensus story says AI opens the field. Protection moves from mechanism to data, data carries no twenty-year clock, and the challengers get in.

The arithmetic runs the other way. Data accrues to installed base. The incumbent generates procedure volume no entrant can approach, and the gap compounds annually, so a protection regime built on data favors whoever already holds the systems. AI is likelier to widen the incumbent's position than to democratize entry, and the effect is stronger the longer that collection has been running.

Two things bound this claim. The first is volume. An executive of a major AI compute supplier, speaking in a published interview at the Society of Robotic Surgery meeting in July 2026, argued that the field holds only a small fraction of the data a robust model would require. If that is right, the incumbent's advantage is blunted rather than decisive, because nobody yet has enough. The remark did not address synthetic data or application-specific models, and the speaker's employer sells the compute that training consumes, so the claim carries an interest.

The second bound is synthetic data. If it proves adequate for a given application, model quality decouples from installed base entirely. Neither bound is settled, and both are worth watching more closely than the patent cliff.

One narrow opening exists. A de novo on an AI-enabled system creates the category predicate, and an authorized predetermined change control plan lets the cleared model be updated within a pre-agreed envelope without a new submission. The combination compounds for whoever arrives first, and it is the single place where the technology favors an entrant.

A narrow door, and it closes behind the first company through it.

None of it matters if the first gate failed

Vicarious Surgical assembled 39 patent families, held 63 active grants with 78 applications pending, and drew 44 citations from Intuitive, 37 from Globus Medical and 13 from Cilag, which is Johnson & Johnson's surgical subsidiary.

The incumbents were reading those patents and designing around them. On any screen built around patent quality, that company passes.

It dissolved in July 2026 having never recorded revenue.

Freedom to sell something nobody needs is freedom of no commercial consequence. The second gate defends a position. It does not create one.

Part IV: Gate 3, The Demand Equation

How many systems will the market take, at what price, with what feature set?

Most people left introductory economics believing it was abstract. Supply, demand, elasticity, curves on a whiteboard with no particular application. This is the application, and this is how it works.

People face nearly infinite choices against finite time and money, and they have to allocate both. Every purchase forgoes something else. The demand equation asks what quantity gets taken at a given price for a specified product, and the enterprise rests on the answer.

The word that does the damage

Entrepreneurs talk about **need** constantly.

Sometimes the **need** is large, sometimes trivial, and the word does not distinguish between them. I can imagine the argument inside a company building a system to do what the incumbent does, differently: here is an approach the market does not yet have. I have heard management at other robotics companies make that argument for years.

Feature differentiators, presented as **needs**, and a dangerous way to approach a market for anyone hoping to build something substantial.

The confirmation problem compounds it. Ask a surgeon whether she needs a capability and she will usually say yes. Agreeing costs nothing, and nobody wants to discourage someone trying to build something. Nobody wants to offend an aspiring unicorn. The people a founder asks are the people least likely to say no, and the founder mistakes politeness for evidence.

The countermeasure is revealed preference over stated preference. What did you do with your last three of these patients? What did it cost? What is in your capital request this year? Those are observations. "Would you use this" is not data.

Whose price

In medicine, the demand equation has four parties.

The **surgeon** specifies the technology and is close to price-insensitive. The **hospital** holds the capital budget and is highly price-sensitive. The **payer**, usually an insurer or a government program, decides whether the procedure carries margin at all, and can foreclose a market without ever being consulted. The **patient** enters as a buyer under specific conditions. A founder who has modeled one has modeled none.

Four coupled constraints, and four different answers to the question of whose P appears in the equation.

When the patient is a buyer

The common formulation says orthopedics is patient-driven and soft tissue is hospital-driven. Prostate surgery is the counterexample.

The real test has three conditions. Is the procedure **elective and scheduled**? Can the patient **choose where the procedure is performed**? Does a poor outcome carry a **vivid, permanent, personal** cost?

Prostatectomy clears all three, and so does elective joint replacement. Emergency or urgent general surgery clears none. Specialty is not the variable.

Where the patient is a buyer, perception moves volume, and perception is not capability.

Units, not enthusiasm

The output of this gate is a schedule. Quantity at price, for a configuration specified tightly enough to build.

The last condition carries more weight than it appears to. A demand estimate for "a surgical robot" is worthless. A demand estimate for a defined system, with a defined feature set, at a defined price, against a defined procedure, is a number you can act on, and it tells you which features carry value and which carry none. The difference between a product and a wish list.

Formal instrumentation exists. Discrete choice experiments and conjoint methods, which present buyers with trade-offs rather than questions, were built to decompose willingness to pay by attribute. How to run them, and how to bound a curve when the procedure does not yet exist, belongs to the companion paper.

Valuation turns on this. Will we reach the market? What will it cost to get there with this precise product? How many will we sell, and at what price? A demanding set of questions, and nobody answers all of them well. Approaching them honestly still beats falling in love with an idea and calling the feeling a market.

The friction nobody prices

A demand estimate built from clinical value and price will still miss, because the purchase runs through an operating room.

Add forty-five minutes to room turnover and the hospital loses a case a day. Require a dedicated scrub technician and somebody has to hire one. Demand a separate sterilization pathway and the central supply department acquires a vote nobody asked it for.

None of that shows up in a clinical value proposition. All of it shows up in the purchase decision. Surgeon enthusiasm rarely survives contact with a theatre manager who has to find the time.

Count the friction in minutes, in headcount and in process steps, then price it. It belongs inside the demand equation rather than in a footnote about implementation.

Why volume is a gate and not a detail

A company can clear the first gate completely and have no business: real unmet capability, no substitute of any kind, a working machine, and thirty units of annual demand, because the procedure is rare, or the anatomy narrow, or the qualified operator population small.

Thirty units supports nothing. Not the regulatory program, not the clinical program, not the field service organization, the manufacturing line, or the sales force needed to place and maintain systems. The capability is real and the enterprise is not.

The case comes up often and gets misdiagnosed as a first-gate failure, which it is not. The clinical claim was true. The market was small. Different errors, different lessons, and collapsing them teaches founders the wrong one.

Volume is also what gives the next gate meaning. Recurring revenue per system multiplies installed base, so a rich consumable stream on one placed system comes to nothing at all. **Gate 4 is arithmetic performed on the output of Gate 3.**

The second demand equation

Who buys the company is also a demand question, and it is partly answerable at founding.

Partly, because what you can do at founding is reason about who might have a strategic imperative to acquire. Real work, available on day one, and it produces a direction rather than an answer.

The arithmetic that governs it

An acquirer's willingness to pay runs roughly to the new profit pool the asset opens, less the margin it destroys on the acquirer's own installed base, less integration cost.

For a product that substitutes for something the acquirer already sells, the middle term is large, and it scales with quality. A better version of what they already sell erodes their base faster than a worse one, so **the number can go negative on an excellent product.**

One plus one comes to less than two here, always. Arithmetic, not preference. Any sensible buyer discounts a similar product entering a market it already serves, because the acquisition cannibalizes the incumbent line.

The practical consequence: your buyer is almost never a company holding the same product, and rarely one holding a similar product with different features. It is a company that lacks the capability entirely and cannot get it anywhere else.

The corollary founders get backwards

If the acquirer's value is what *they* can sell through *their* channel at *their* margin, then your sales force, your distribution agreements and your marketing program are not assets in a strategic transaction. They are redundant cost, eliminated after close.

Founders build them out believing they raise acquisition value. Against a strategic buyer they frequently destroy it, or at best spend capital on something the buyer will discard. The company is not valued on what it can sell standalone with a fledgling commercial organization. It is valued on what the acquirer can sell, and what the acquirer can earn.

Forecasting the gap

The harder problem is that the buyer has to be projected forward, to the moment your company offers a transaction that makes them money, which is rarely the moment you founded. Markets change, budgets change, incentives change.

You are not forecasting the acquirer's current portfolio. You are forecasting their **gap**: what they will be missing in five to seven years, given their own roadmap and their principal competitor's trajectory. Skate to where the puck is going.

The gap at Johnson & Johnson that made Auris valuable in 2019 did not exist in 2007. Its shape was foreseeable, because the strategics were visibly building programs that duplicated or directly competed with the market leader, and a company doing that eventually needs the thing it is not building.

The limit

The pool is unstable. Obvious buyers vanish over strategy, timing, leadership change or a competing acquisition, and buyers nobody modeled appear. Any claim that the acquirer can be named at founding is storytelling after the fact.

The defensible claim is narrower and still worth a great deal. The reasoning is available at founding, it eliminates whole categories of buyer on arithmetic, and it tells you which capability gap to aim at.

The signal in the other direction

When a strategic investor sits inside the capitalization table for years with full visibility and declines to acquire, the party best positioned to render a demand verdict has rendered one. Becton Dickinson entered the 2021 Vicarious placement, sold in the 2023 secondary, and exited holding nothing.

The limit of the equation

For a new category no observable curve exists.

In 2007 there was no revealed preference for peripheral lung access, because the procedure did not exist. There was nothing to observe, and no purchase history to reason from. The curve can be

bounded from adjacent observables: nodule incidence, existing biopsy volumes, diagnostic yield failure rates, downstream treatment costs, the size of the population currently receiving nothing. It cannot be measured.

So this gate produces a bounded estimate, not a measurement.

The honest statement, and it does not weaken the requirement. **What stays uncertain after diligent bounding is small next to what stays uncertain when the bounding never happens.** Failure in this field is almost never a demand estimate that proved wrong. It is a demand estimate nobody constructed, replaced by a category-level market size and a run of agreeable conversations.

Part V: The Fourth Gate, and Why the Order Is the Argument

It’s hard to get rich selling capital equipment

Selling capital equipment is often a poor business on its own terms. Eighteen-month sales cycles, revenue arriving in blocks and then stopping, price pressure at every renewal, and a service obligation that outlives the margin on the original sale.

The durable businesses in this field were never robot businesses. They were consumable businesses with a robot attached. End effectors wear out. They are disposable or reposable, never infinitely reusable, and every procedure consumes them.

Does each installed system produce recurring high-margin revenue?

The answers are the least glamorous objects in the room. Instruments with limited use lives. A single-use bronchoscope and sheath, replaced per procedure, with the manufacturer holding clearance to perform the one permitted reprocessing cycle itself. An implant only this system will accept.

The strongest form is not a razor and blade at all

The strongest form is a robot built to sell a high-margin product the company already manufactures.

Stryker did not buy Mako to enter the robot business. It bought a machine that, once installed in an operating room, obliges the surgeon to use Stryker knee implants. The capital equipment functions as a customer acquisition cost that happens to sit on the balance sheet as a product.

Johnson & Johnson's VELYS does the same job defensively for the ATTUNE implant line, and OTTAVA does it for Ethicon's staplers and energy devices. Neither robot had to clear the first gate on its own account, because the annuity was already owned.

A real strategy, available to perhaps six companies on earth. Part IX takes it up properly, since it bounds this framework rather than applying it.

Software is the best instrument the field has produced

A subscription on an installed base carries no cost of goods. Nothing to manufacture, sterilize, ship or stock. No unit that has to exist before it can be sold.

It converts capital equipment into an annuity, and the major players are visibly building toward it.

The current enthusiasm for AI in this sector therefore reads better as a Gate 4 strategy than a Gate 1 one. The value is not that the algorithm makes the surgery possible. The value is that the algorithm can be billed monthly.

The arithmetic that makes this a gate

Recurring revenue per system multiplies installed base.

Multiply a superb consumable margin by one placed system and you have nothing. Multiply it by four thousand and you have Intuitive Surgical.

So this gate performs arithmetic on the output of Gate 3. It cannot be assessed on its own, it cannot rescue a small market, and a company that has not cleared the first three will never place enough systems for consumable economics to matter at all.

Gate 4 is moot without Gate 3. Gate 3 is moot without Gates 1 and 2.

What the arithmetic has to clear

Margin is not the only thing at stake. This gate determines whether the company can return a fund.

Venture portfolios depend on the winners carrying everything else, which means an outcome several times the capital deployed. Run that backwards. A capital sale with no recurring revenue has to place an extraordinary number of systems to reach the number, against a sales cycle measured in quarters and an installed base measured in hundreds. The arithmetic rarely closes.

Attach a consumable and the same installed base compounds. A system placed once keeps paying, so terminal value stops depending on how many machines you sell this year and starts depending on how many you have sold to date.

Boards ask about the consumable long before they ask about the robot for that reason. They are not asking about margin. They are asking whether the exit arithmetic works at all.

Why the order is the argument

State these four as a list and they are unremarkable. Substitutes, freedom to operate, meaning the right to sell without being sued out of the market, demand, and recurring revenue. Any experienced investor in the field would name all four.

State them as ordered gates and they become falsifiable.

The framework predicts that a company failing the first gate will not be rescued by an excellent answer to any of the others, whatever its engineering quality, capital raised, patent strength or investor pedigree. It predicts that diligence performed on gates two through four, before the first is settled, is diligence on a question the company has not yet earned.

Both predictions can be checked against the record, and Parts VI through VIII do that.

Where AI creates value

Artificial intelligence has its greatest impact on Gates 2 and 4, both of which become irrelevant if Gate 1 fails.

In other words: **AI makes a company that cleared the first gate stronger, and does nothing whatever for a company that did not.**

Perception is the exception. Where the limiting factor is seeing rather than manipulating, the machine can do what the hand cannot: distinguish tissue the eye reads as identical, find a margin nobody can see, hold registration when anatomy moves and the preoperative scan stops being true.

That is a first-gate claim and it will survive scrutiny.

Workflow analytics, scheduling optimization and skill assessment will not. They are Gate 4 products wearing Gate 1 language, and the distinction will separate the next cohort of survivors from the next cohort of case studies.

Part VI: The Gates Applied

Seven companies. Same four questions, asked in the same order.

These seven were not chosen at random. They were selected for the depth and verifiability of their public record: filings, clearance documents and patent data sufficient to code all four gates without guessing. That rule has a consequence worth stating. Data-richness correlates with public listing, so the set over-represents companies that raised public capital and under-represents private firms, non-US firms and divisions of larger companies. Whether that biases the result in either direction has not been tested. Companies excluded on that basis, among them Hansen Medical, Titan Medical, CMR Surgical, Distalmotion and the Medtronic Hugo program, appear in the text where relevant and are not coded in the matrix.

The table below is blunt on purpose. Every cell is a judgment and several are arguable, which is the point: a framework nobody can argue with is a framework nobody can check. The walkthroughs take up the four cases where the answer is not obvious.

Read the rows for the pattern rather than the verdicts.

Intuitive, and where the pass came from

Easy to say the market leader cleared the first gate because patients wanted the robot. The common reading, and it is wrong. Getting it wrong has cost founders a great deal of money.

Laparoscopic radical prostatectomy existed before da Vinci and delivered a comparable clinical result at far lower patient burden than open surgery. Almost nobody could perform it. The pass came on reproducibility: the machine made a punishing operation available from ordinary surgeons at scale.

Patient perception is a separate fact and it belongs at Gate 3. Prostatectomy is elective and scheduled, the patient chooses the site of care, and the downside is vivid and permanent, which is why perception moved volume so fast. It accelerated an adoption curve. It did not create the answer underneath.

Mako repeats the structure in orthopedics. A surgeon's hand can replace a knee, so clinical result fails. Reproducibility passes: consistent alignment and placement from a general population of surgeons rather than from the best ones, with patient perception again doing demand work at Gate 3.

Hold onto the distinction. A company with only the second half has a marketing plan attached to a substitute.

The Auris outcome is shown two ways, headline and paid, because the difference between them is the subject of Part 1 of this series.

Medrobotics: the case that should have worked

This row matters most, because it falsifies the simple version of the framework.

The Flex Robotic System cleared FDA in 2015 and held CE marking, the European equivalent. It did what rigid instruments could not: articulate around corners to reach anatomy in the pharynx a straight scope cannot see. On the first gate, on access, a genuine pass.

It reached roughly thirty hospitals and filed for Chapter 7 in January 2022 against more than $135 million in asserted claims.

The third gate is the one it failed. Transoral robotic procedures in the United States numbered 994 in 2016, and the literature defines a high-volume center as one performing ten a year. A real clinical need, and nowhere near the volume needed to carry a regulatory program, a manufacturing line, a field service organization and a sales force placing systems one hospital at a time.

It also faced a substitute question it never fully answered, since da Vinci had held transoral clearance since December 2009, six years before Flex arrived.

Real capability. Real clearance. Real hospitals. No market.

A three-gate framework has nowhere to put that case. Demand is a gate here for that reason, and not a footnote. Under the older formulation the only available classification was first-gate failure, and that is false. The clinical claim was true. The market was small. Different errors, different lessons.

SS Innovations: the gates clearing in different countries

This company trades, with an FDA submission filed in December 2025 and no outcome announced as of this writing. Nothing here forecasts its result, and nothing should.

What interests me is that its answers to the first gate are split by geography.

On access it holds a real pass. Cardiac telesurgery across roughly twenty thousand kilometres of network distance at latency near three hundred milliseconds is not something any substitute delivers for a patient with no surgeon in range.

In the United States the proposition is a multi-specialty robot at a lower price, competing against installed da Vinci systems and now Hugo. A substitute-rich position, and the company's own filing says so: it chose a 510(k) over a de novo, which is a written assertion that a predicate exists.

The company has disclosed something about that filing worth reading closely. The December 2025 submission covers the surgical system across general, urological, colorectal, gynecological and cardiac indications, and the company has said separately that telesurgery authorization would follow FDA and CE marking of the system itself.

On the company's own account, then, the capability that clears the first gate is not the one currently before the agency. What is on file is the robot, in anatomy already occupied, on an assertion of equivalence to what is already there.

So the gate clears where the company cannot easily monetize it and fails where it could. The demand equation and the financing problem that follows come up in Part IX.

VELYS: the row that should not work and does

Run it through and it fails comprehensively. A surgeon's hand replaces a knee. It entered a market already holding Mako, ROSA and CORI, so it arrived fourth to a party rather than first.

By this framework it should be dead. It is cleared in twenty markets with more than fifty-five thousand procedures performed.

The explanation is that VELYS is not a company at all. It is a product line inside an organization that already owned the fourth gate, and the robot's job was never to create value. Its job was to stop the loss of value already captured, by making sure a surgeon did not have to leave DePuy to get a robot.

A real limit on the framework rather than an application of it, and Part IX treats it properly.

The Gates Applied

Seven companies, four questions, asked in the same order.

Read the rows for the pattern rather than the verdicts. Reasons are in the walkthroughs.

	GATE 1 SUBSTITUTES	GATE 2 RIGHT TO	GATE 3 DEMAND	GATE 4 CAPTURE	OUTCOME
Intuitive da Vinci	PASS	PASS	PASS	PASS	Category owner
Auris Monarch	PASS	PASS	PASS	PASS	$5.75B headline $4.2B paid
Mako Robotic knee	PASS	PASS	PASS	PASS	$1.65B, 2013
Medrobotics Flex	PASS	PART	FAIL	—	Chapter 7, 2022
Vicarious Single-port	FAIL	PASS	—	—	Dissolved, 2026
SS Innovations SSi Mantra	PART	PART	PART	—	Trading; filing pending
VELYS Robotic knee	FAIL	n/a	n/a	PASS	Thriving

Rows that bound the framework rather than illustrate it. Medrobotics cleared Gate 1 and failed anyway. VELYS fails Gate 1 and thrives inside a parent that already owns Gate 4.

Across seven companies, no answer to Gates 2, 3 or 4 ever rescued a failure at Gate 1.

Vicarious held a portfolio the market leader was designing around. Medrobotics held a real capability with a market too small to carry it. Failing the first gate is decisive.

Figure 3. Seven companies, four questions. No answer downstream ever rescued a failure at Gate 1.

What the matrix is claiming

Not that these seven verdicts are beyond argument. Several are arguable and the walkthroughs argue with them.

The claim is narrower and harder to dismiss. **Across seven companies, no answer at gates two, three or four ever rescued a failure at gate one. And no company that cleared all four failed for a reason this framework is designed to explain.**

Both anomalies clarify the framework rather than weaken it. Medrobotics shows that clearing the first gate is necessary and nowhere near sufficient. VELYS shows the framework governs standalone companies, not products launched inside firms that already own the recurring revenue.

Limits, both of them, and a framework that states its limits is a framework you can use.

Part VII: The Case

Vicarious Surgical, incorporated in 2014 to build a single-port surgical robot, ceased operations on 21 July 2026. It consumed $253.4 million, the accumulated deficit reported in its final quarterly filing. It reported no revenue in any fiscal year.

Everything that follows comes from public filings. Companies are named; individuals are not, because the failures here are structural and no person needs to be in them.

Gate 2 first, because it passed

The instinct is to look for missing intellectual property. It is not there.

The company assembled 39 patent families, 188 documents, with 63 grants active and 78 applications pending at dissolution. Its most-cited patent carries 91 forward citations and the next carries 84.

More telling is who was citing the portfolio. Intuitive Surgical cited this portfolio 44 times. Globus Medical cited it 37 times. Distalmotion, a European single-port entrant, 17 times. Cilag, Johnson & Johnson's surgical subsidiary, 13 times.

The incumbents were reading these patents and designing around them.

Any diligence process organized around patent quality would have cleared this company, and by the look of it several did. The portfolio was real, it was cited, and it did not matter, because the gate it satisfied sits downstream of the gate that failed.

Gate 2 passed. Gate 1 was never cleared. The company is gone.

Gate 1(a): the substitute was there before the company was

The system aimed at abdominal soft tissue procedures, with ventral hernia repair as the announced beachhead.

Laparoscopic ventral hernia repair is routine, high volume and adequately reimbursed, performed thousands of times a day. The human hand, working through conventional laparoscopic instruments, gets the clinical outcome.

No version of the outcome-specified substitute test passes here. The procedure was not impossible, not difficult to access, and not underserved.

Gate 1(b): the differentiation, and how it eroded

The differentiation was architectural. A single port, two human-like articulated arms working inside the body, a virtual reality interface.

Intuitive received initial clearance for its own single-port platform in April 2014, refined it, and cleared it again for urologic procedures in May 2018, saying at the time that it intended to pursue transoral, transanal and extraperitoneal applications and broaden the platform over time.

A published roadmap from the market leader, pointed at the architecture Vicarious was building toward. Single-port clearances for inguinal hernia repair, cholecystectomy and appendectomy followed in December 2025, and Medtronic's Hugo reached its own first US clearance in the same period.

So the architectural claim carried a shelf life visible from 2018 onward. But be precise about what this does and does not establish. It does not mark the moment a substitute appeared, because on the outcome that mattered, repairing a ventral hernia, the substitute had been there the whole time.

It establishes something narrower and still useful: even the differentiation that was real ran on a public clock, held by an incumbent that had announced where it was going.

The founding artifact

The sequencing inversion here is not inferred from strategy documents or interviews. It is visible in the first thing the company created.

The earliest patent filing is dated 5 May 2014, the year of incorporation, titled "Virtual reality surgical device." The four next-most-cited filings are a virtual reality surgical device, a virtual reality wrist assembly, a virtual reality surgical camera system, and a further virtual reality surgical device.

An architecture, not an indication.

A company beginning with a clinical problem produces a first filing describing access to anatomy. A company beginning with a machine produces a first filing describing the machine. The record is unambiguous about which happened here, and the beachhead indication arrives years later, chosen to fit a platform that already existed.

The inversion described in Part II, documented at the level of a filing date.

What the financing record shows

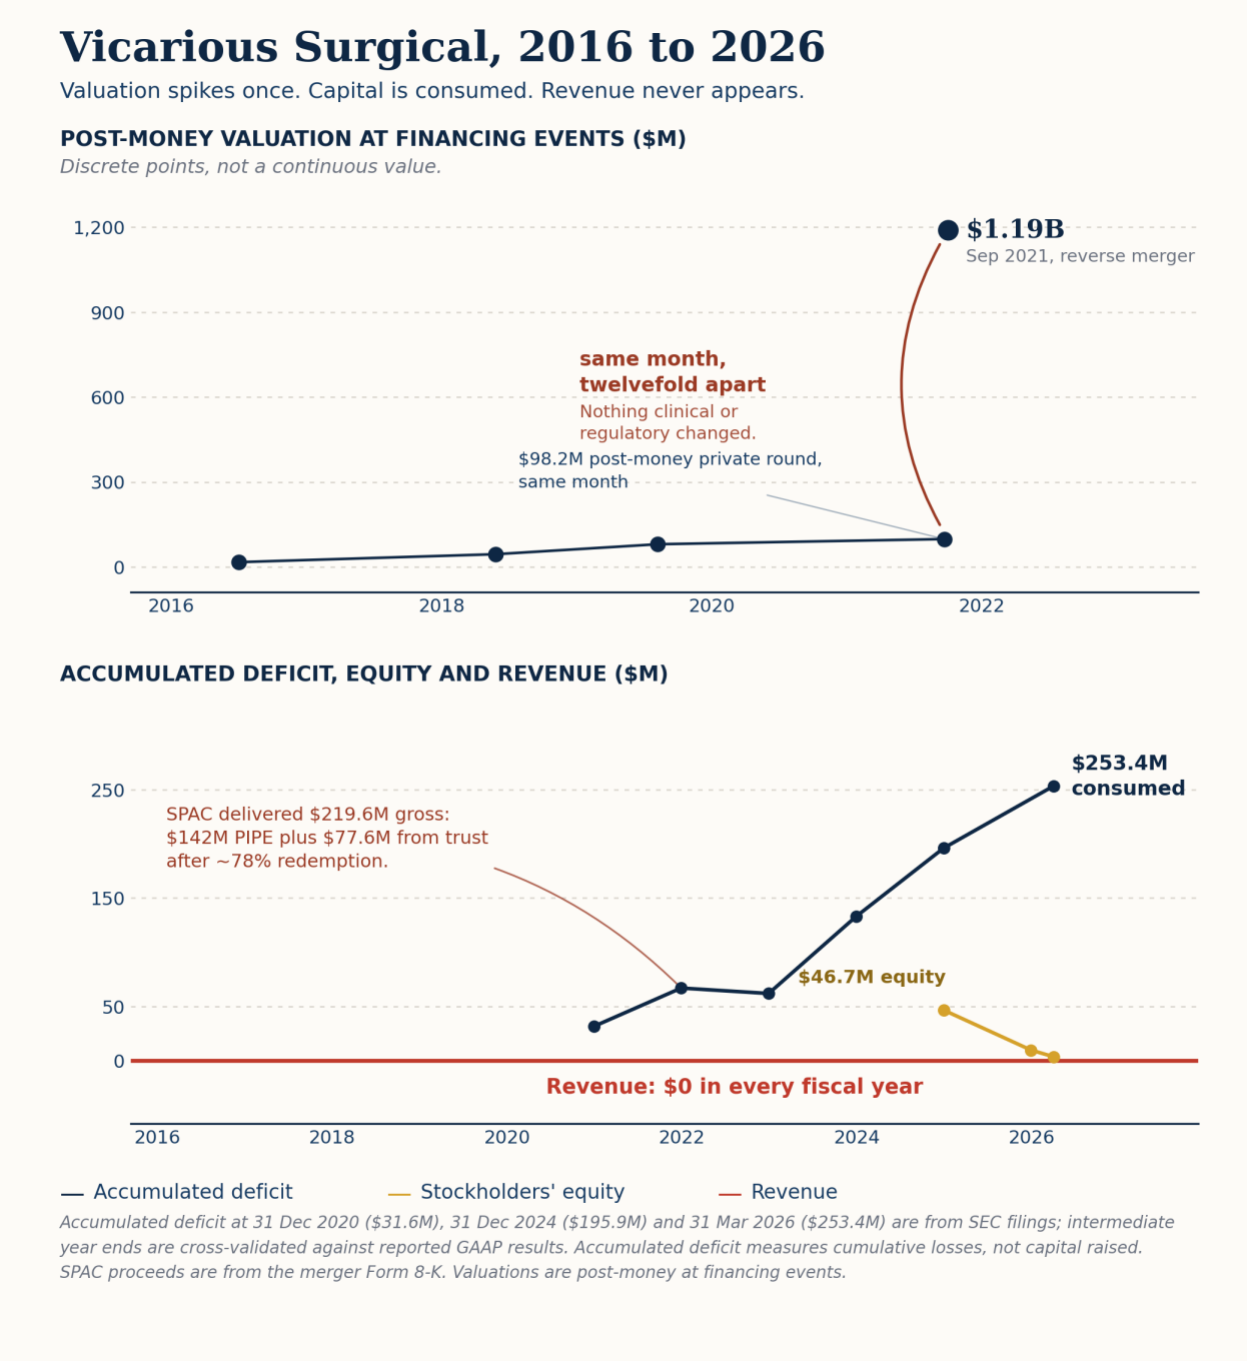

Figure 4. Vicarious Surgical, 2016 to 2026. Valuation spikes once; capital is consumed; revenue never appears.

The private rounds were modest and progressively priced. A 2016 Series A raised $4.19 million at a $16.74 million post-money valuation. A 2018 round raised $16.76 million at $44.76 million. A 2019 round raised $10.00 million at $80.00 million. A September 2021 round raised $13.20 million at $98.20 million.

In the same month as that last round, a reverse merger with a SPAC set a $744 million pre-money and a post-money valuation estimated at $1.19 billion. The trust held roughly $345 million. After redemptions it delivered $77.6 million, alongside a $142 million private placement, for $219.6 million of gross proceeds.

Twelvefold in a single month. Nothing clinical, regulatory or competitive changed in that month. The venue changed.

What follows is a company spending it down. Operating losses of $66.6 million in fiscal 2024, $51.0 million in fiscal 2025, and $42.7 million in the twelve months ended March 2026. Total equity falling from $46.7 million to $9.8 million to $3.6 million across the same periods. Investing activities contributing $32.1 million and $32.6 million of cash in the last two periods, which is the signature of a securities portfolio being liquidated to fund operations rather than an investment program.

A follow-on share sale raised $45 million in 2023. A final private placement in October 2025 raised roughly $5.87 million, by which point the shares had moved from a national exchange to over-the-counter quotation.

Twelve years. No revenue in any year. No clearance in the record.

The strategic that was already inside

Becton Dickinson participated in the 2021 private placement, sold in the 2023 secondary, and exited at dissolution holding nothing.

The quiet data point in the file. The natural acquirer did not have to be found, courted or educated. It sat in the capitalization table with board-adjacent visibility for years and did not buy the company.

When a strategic investor with full access declines to acquire across a five-year horizon, the party best placed to render a verdict on Gate 1 has likely rendered one.

What did not cause this

The trade explanations are all true and all downstream.

The financing structure. The shell merger, or de-SPAC, put a pre-clinical hardware program into a venue that prices narrative and performs no device diligence. It determined the shape of the failure. It did not determine that there would be one.

The burn rate. Tens of millions a year against zero revenue is fatal only for a company that cannot reach a market. Every successful company in this field burned comparable sums for comparable periods.

Interest rates. Capital conditions tightened for everyone. A company with a cleared Gate 1 answer raises in a hard market at a bad price. A company without one raises only while the market is still willing to price narrative.

Competition. Competitors did advance. Not a shock arriving from outside, but the Gate 1(b) answer that had been public since 2018.

Each explains the timing and the manner. None explains the outcome. Remove any one of them and the company still has no answer to what operation becomes possible.

The knowability gradient

The useful question is not why it failed but when the failure became assessable.

At founding, 2014. The substitute test fails on the announced indication. Laparoscopic ventral hernia repair was routine, high volume and adequately reimbursed, and multiport robotic hernia repair already existed. No forecast required. Reading the target indication and asking what already achieved that outcome would have done it.

Reinforced through 2018. The incumbent's single-port platform cleared for urology and published its intent to broaden. The architectural differentiation now carried a visible clock, held by somebody else.

Beyond argument by September 2021. Seven years in. No revenue in any year. No clearance. No substitute-free indication. Priced at roughly $1.19 billion in a reverse merger, in the same month a private round priced the same company at $98.2 million.

Nothing after that date was a bet on the first gate. It was capital underwriting execution against a question answered at incorporation.

What the case establishes

Not that surgical robotics is hard. It is.

The case establishes that **Gate 1 is assessable in advance, from public information, at datable moments, and that here it was not assessed.** The patents were counted. The engineering was evaluated. The market was sized at category level. The one question that decided the outcome, whether any substitute achieved the same clinical result, was available to anyone who asked it and appears to have governed no financing decision after 2018.

The failure worth studying, and not a failure of intelligence. Every party here was sophisticated. It is a failure of sequence.

Part VIII: The Falsification Test

A framework that only explains winners explains nothing. Any consultant can assemble three successes and derive the principle they share.

So this section does the opposite. Two cases that should break the argument, and what survives them.

Case one: a company that cleared the gate and died

Medrobotics appeared in Part VI. Here is why it belongs in a section about falsification rather than in a table.

Every element the framework says matters was present. A capability rigid instruments did not have. FDA clearance in 2015. CE marking. Roughly thirty hospitals, meaning real surgeons chose it with real capital budgets. And in January 2022, Chapter 7, against more than $135 million in asserted claims.

If the argument were "clear the first gate and you will be fine," this case ends it.

Medrobotics proves that **necessary is not sufficient**, and the paper is stronger for putting it in the text rather than waiting for a reviewer to raise it.

Volume is what it failed, and the numbers are not close.

Transoral robotic surgery, performed through the mouth rather than through an incision in the neck, ran to 994 cases in the United States in 2016, up from 363 in 2010. Across seven years fewer than five thousand patients received it, out of 73,661 diagnosed with oropharyngeal cancer over the same period; most were treated without surgery at all. The literature defines a high-volume center as one performing ten or more cases a year.

Thirty hospitals at that rate is three hundred procedures.

No business plan repairs a denominator that size. You cannot market your way out of it, and you cannot shrink a field service organization below the number of sites that exist.

The lesson is not that Medrobotics was a bad company. It is that a true clinical claim and a viable business are different findings, and the first feels so much like the second that founders and investors stop looking.

Case two: the same pattern, outside surgery entirely

Figure 5. Two clocks. Regulatory time is a moat if you are through it and a trap if you are not.

A framework built from seven surgical robotics companies could be nothing more than a description of those seven. The only way to find out is to run it somewhere else. So take a field with no robots in it.

Motion preservation of the spine was the most heavily capitalized story in orthopedics through the 2000s. The premise was excellent. Spinal fusion eliminates motion at the treated level and transfers load to the adjacent segments. An artificial disc preserves motion and, in theory, protects the neighboring anatomy. Billions went in.

The FDA approved the Charité artificial disc, from DePuy Spine, on 26 October 2004. It approved ProDisc-L, from Synthes Spine, on 14 August 2006.

Now look at what happened, and when.

On 16 May 2006, the Centers for Medicare and Medicaid Services issued a national coverage determination on lumbar artificial disc replacement, the ruling that decides whether Medicare pays for a procedure at all. It found the procedure not reasonable and necessary for beneficiaries over sixty and issued national noncoverage for that population, leaving everyone sixty and under to local contractors. The request had come from outside the industry, from a physician asking for noncoverage. CMS opened that determination for reconsideration in November 2006 and released a decision memo on 14 August 2007 which widened the ruling, replacing noncoverage of one named device with noncoverage of the procedure itself.

Read the sequence. The first device was approved in late 2004. Inside nineteen months the largest payer in the United States had opened a formal position on whether it would pay for the procedure at all. The second device received FDA approval three months after the coverage determination issued.

Every follower in that category was, at that point, running an expensive multi-year regulatory program aimed at a market whose reimbursement was being decided by a party none of them controlled and most of them had not modeled.

The rest is record. Charité was withdrawn in 2012. The next lumbar total disc replacement approved by FDA was activL, in June 2015, nine years after ProDisc-L. A category that had absorbed billions produced no new approved lumbar device for close to a decade.

What the spine case demonstrates

Two things, and neither is the usual lesson. The spine case is normally told as a story about devices that failed clinically. That is not what happened here.

The payer is a customer, and the payer can foreclose a market without ever being consulted. The four-customer problem from Gate 3 in its most consequential form. The surgeons wanted it. The patients wanted it. The hospitals bought it. The clinical rationale was sound and remains sound. None of it mattered, because the party deciding whether the procedure carried margin was working to its own timetable and its own evidentiary standard.

A demand equation that models surgeon enthusiasm and hospital budgets and stops there is a survey.

And the barrier that closed the category to entrants was time, not patents. A premarket approval, the FDA's most demanding pathway, run with an investigational device exemption and multi-year follow-up, meant a follower's answer had to hold five to seven years out. It did not. By the time a fast follower could ship, someone else had settled the question, and settled it badly.

The point from Gate 1 with dates attached: you do not run the test against today's field, you run it against the field on the day you can ship. The longer your regulatory pathway, the further into the future your answer has to survive.

One structural difference

Spine motion preservation had a single target. The disc. One anatomy, one race, and when it was run it was over for everyone.

Surgical robotics has dozens of anatomies. Lung, prostate, knee, colon, hernia, pharynx, heart, spine, and the ones nobody has taken yet.

The difference matters enormously to a founder. **Foreclosure in robotics runs per anatomy, not across the field.** The entrants crowding into abdominal soft tissue are running the spine race, competing for a target already claimed. The ones selecting anatomy nobody has taken are not in that race at all.

So the gold rush framing is half right. Capital does flood in after the category-defining exits, and most of it is wasted. But it is wasted on the specific anatomies where the answer is already known, by companies that never asked which race they had entered.

What survives

Both cases were chosen to break the framework and neither does. What they do is bound it.

Clearing the first gate is necessary and not sufficient. Volume can kill a company holding a true clinical claim, and a payer can close a market everyone else in the chain wanted open.

Neither is a reason to skip the first gate. Both are reasons to run all four.

Part IX: The Limits of the Framework

A framework that explains everything explains nothing. These four gates are necessary conditions for an independent company to build durable value. They are not sufficient, and three classes of company fall outside them entirely.

Stating the limits is not a hedge. It separates a framework from a slogan, and it is where the practical guidance lives.

Limit 1: the incumbent that already owns Gate 4

The gates govern standalone companies that must clear all four to survive. They say nothing about a product line inside a company that already holds the annuity.

Consider VELYS, the DePuy Synthes robotic knee platform cleared in 2021 for use with the ATTUNE total knee system. Run it through the gates and it fails comprehensively. A surgeon's hand performs a knee replacement, so the clinical result dimension fails. It entered a market already holding three established orthopedic robots, so it arrived fourth rather than first.

By the framework it should be dead. It is instead cleared in twenty markets with more than fifty-five thousand procedures performed and a unicompartmental indication granted in 2024.

The explanation is that the robot's job was never to create value. Its job was to stop the loss of value already captured. When Stryker uses Mako to pull its own implants through, the defensive response is a robot of your own, and it does not have to be better. It has to exist, so the surgeon does not have to leave to get one.

The tell is the origin. Orthotaxy, the French developer it came from, was acquired in 2018 for an undisclosed sum and had no standalone future. It had value only inside a company with an implant franchise to defend.

The pattern is not confined to orthopedics. OTTAVA performs the identical function for Ethicon's stapler and energy franchise. Both robots are Gate 4 instruments, deployed by a company that already owned the consumable. Neither had to clear Gate 1 on its own account.

The founder-facing consequence is the part that matters.

If your first-gate answer is weak, your company is not worthless. It is worth something specific: whatever an incumbent will pay for a defensive instrument. A real outcome, and not the one in your model. An undisclosed sum for a technology that could not stand alone, rather than a multibillion-dollar acquisition of a company that could.

Founders routinely mistake the first for a smaller version of the second. Different transactions, different buyers, different logic. Knowing which one you are building toward is a Gate 1 question, and it should be answered in year one rather than in the data room.

Limit 2: the capital markets have no stage gate for devices

Whether a company that clears the gates can finance the crossing is a separate question, and the reason is structural rather than circumstantial.

Pharmaceutical development has a shared vocabulary for partial de-risking. Phase 1, Phase 2, Phase 3, with published probability-of-success figures by phase and indication. A readout reprices the asset on a scale every analyst accepts. A company that has sold nothing can still be valued, because the industry agreed long ago on what counts as progress.

Medical devices have no equivalent. No analyst framework reprices a company for completing a design freeze, closing a pre-submission meeting, or finishing a human factors validation study. The clinical process is not standardized across device types, the regulatory pathway varies by predicate and classification, and no published table converts any of it into probability.

So there is no partial credit. A device company is worth its narrative until it is worth a revenue multiple, with almost nothing in between. Public markets price uncertainty poorly under the best conditions, and an unbenchmarked regulatory pathway is a large uncertainty. Compounding it, what public markets want from this sector is proof: shipped systems, revenue, earnings.

And here the framework produces a tension it has to acknowledge.

A strong answer at Gate 1 means no substitute exists. No substitute frequently means no predicate device. No predicate means a de novo pathway, which is the pathway a generalist analyst cannot benchmark against anything.

The better your answer to the first gate, the less legible you are to public capital.

Which suggests, painfully, that companies with the strongest first-gate answers should stay private and exit to a strategic, and that public equity is structurally mismatched to this asset class at exactly the stage where the capital requirement peaks. It also explains a pre-clinical hardware program arriving on a public exchange through a reverse merger better than any account of individual judgment. The venue was not chosen because it suited. It was chosen because it was open.

Hansen Medical is the illustration, and its financing history tells the story without commentary.

It went public in November 2006, raising $78.3 million net. It obtained CE marking that quarter, FDA clearance in the second quarter of 2007, and recorded its first product revenue immediately after. A device company with cleared products and paying customers.

Then watch what the market charged it to raise money. April 2008: three million shares for $39.5 million net. April 2009: 11.7 million shares for $35.3 million. April 2010: 16.1 million shares for $29.8 million. Progressively more equity for progressively less cash.

In April 2016 it sold to Auris for $4.00 a share, roughly $80 million in total equity value, at a 40 percent premium to the previous close. Ten years after an IPO that raised $78.3 million, the whole company changed hands for about the same amount.

Clearance was never the problem. The public market had no way to price what it was looking at, and repriced it downward every year until there was nothing left to reprice.

Everything else can be right. The capability real, the demand equation sound, the freedom to operate clean, and the company can still die in the capital markets, public or private. Nothing here predicts which companies survive that, and any claim otherwise would be storytelling.

And the structure holds the mistake in place

A second mechanism operates here.

A company that raises a large public round against an unresolved first gate does not just hold money. It holds obligations. Public reporting, a float, an investor base that priced the story, and often a preference stack, the order in which investors are repaid, determining who sees anything in any outcome short of the plan.

Those obligations all point forward. Spending goes to commercial overhead, sales hires and market development, because that is what the story promised and what the market watches for. Reopening the first gate remains available in theory and close to impossible in practice.

The capital structure does not cause the failure at Gate 1. It prevents the correction.

Limit 3: timing and acquirer appetite

The gates explain why a company becomes worth acquiring. They do not explain why it was acquired, at that price, in that quarter.

An acquirer's strategic need, the state of its internal programs, the pressure it feels from a dominant competitor, the availability of an alternative: all of it determines outcome, and none of it is a property of the target. A company that clears all four gates in a year when no strategic buyer needs the category will wait, and waiting consumes the capital that clearing the gates required.

Not a defect in the framework but the boundary of what any framework built on company characteristics can reach.

Limit 4: the gates say nothing about whether you can execute

Four gates assess whether the opportunity is real. None asks whether this company can take it.

Manufacturing at quality. A regulatory function. Clinical development. Training and proctoring. Field service. Distribution. Reimbursement competence. Hospital integration. Supply chain. And the organizational learning that makes all of it repeatable rather than heroic.

Teece's work on profiting from innovation describes innovators who created important technology and lost the value to whoever held the complementary assets. Pisano's describes capabilities accumulated across years rather than purchased in a quarter.

Both are real, and neither is a gate. They are not binary, they do not sequence, and they belong to a company rather than to an opportunity. Conflating the two is why "strong team, no market" and "real market, weak team" are both familiar failures with entirely different remedies.

The gates tell you whether the prize exists. They will not tell you whether you are the one who can reach it.

What the gates do claim

Only this: **failing Gate 1 is dispositive, and passing all four guarantees nothing.**

The asymmetry is the useful part. A test that predicted success would have to model everything, and could be run only in hindsight. A test that identifies a fatal absence needs one answer, and that answer is available from public records before the money is committed. A weak first-gate answer cannot be rescued by capital, engineering quality, investor pedigree, regulatory persistence or artificial intelligence. A strong one can still be lost to financing, timing or biology.

A narrower claim than most frameworks make, and the one that survives contact with the record.

What would prove this wrong

Seven companies, selected after their outcomes were known, coded by an author with a view. Theory building, not proof, and no amount of internal consistency repairs the method.

So the claim, stated forward, where it can fail.

If a company that fails Gate 1 by any honest reading goes on to build a durable independent venture-scale business on the strength of its patents, its demand or its consumable stream, the framework is wrong.

If companies now filing 510(k)s into occupied anatomy, asserting equivalence to an installed incumbent, produce independent venture-scale outcomes, the framework is wrong. Several such filings are pending as this is written. The results will arrive on their own schedule and will not require my agreement.

If a field turns up where resolving demand before substitutes reliably produces better outcomes, the ordering claim is wrong, and the ordering claim is the contribution.

What will not falsify it: a company that clears all four gates and fails anyway. This section already concedes financing, timing and biology. A framework has to be told what it does not cover, or every outcome confirms it and it explains nothing.

Part X: Selected Risk, or How Do You Know It Was Not Luck

I have been asked this more than any other question. A single good outcome is indistinguishable from a lucky one, and the person asking is usually right to ask.

You cannot prove an outcome was not luck. Nobody can, in any field, from a sample of one. What you can do is audit the construction, and the construction is visible before the outcome is known.

Luck produces a portfolio of risks you did not choose. Design produces one risk you did.

What the field looked like in 2007

Worth remembering what the evidence said when we started.

Start with what had happened to the market leader, because everyone in the field had watched it.

Intuitive Surgical lost money through 2001 and 2002 on revenue of $51.7 million and $72.0 million. Its first profitable quarter came in the second quarter of 2003, on net income of $0.9 million. Its first profitable year was 2004.

Then look at how the turn happened. Recurring revenue rose from 29 percent of sales in the first quarter of 2003 to 46 percent a year later, up 125 percent, while the installed base climbed from 210 systems to 286 across 2004. A capital equipment company converting into an annuity, and it converted because the machines already placed started getting used.

What started using them was prostate.

Meanwhile the clinical position was close to unanimous. Surgeons said they could do as well with their hands, and the data at the time supported them. Systems sat in hospitals between cases, some of them bought with donor money a salesperson had helped coordinate. Every company in the field was building better instruments for procedures surgeons already performed, and the outcomes literature was not moving.

Two conclusions came out of that, and the company was built on both.

If the honest answer to "can you do this by hand" is yes, no amount of engineering repairs the position. The only durable place to stand is a procedure the hand cannot perform. Better instruments for existing operations is a crowded field with no exit from it.

And what rescued the market leader was not a surgeon's judgment about instruments. It was prostate, and what drove prostate was not the literature.

Between 2003 and 2013, robot-assisted radical prostatectomy went from roughly 1.8 percent of prostatectomies in the United States to 85 percent. The first randomised trial comparing it to open

surgery reported in 2016, found no difference in urinary or sexual function at twelve weeks, and confirmed the same at twenty-four months.

The evidence arrived a decade after the market had been decided.

What decided it was that no man wants to hear his surgeon might leave him incontinent or impotent for the rest of his life. The patient chose the center with the machine, and the hospital followed the patient.

One of those told us what to build. The other told us who decides, a question I had not heard anyone in the field ask out loud.

The inventory

At Auris the exposure was narrowed deliberately, one risk at a time.

Market risk came down by selecting a procedure with no substitute. If the capability was real, demand did not depend on persuading anyone that a better version of an existing procedure was worth switching for.

Competitive risk came down by selecting anatomy the market leader was not in. We were not going to race Intuitive. We were going to pull up beside them with something they had not done, at a moment when the strategics were building programs that largely copied or directly competed with them.

Exit risk came down by naming the acquirer at incorporation. A strategic already holding a duplicate of the leader's platform does not want a second duplicate. It wants the thing it cannot build and cannot buy from anyone else.

Execution risk came down by recruiting an operator with the relevant record, once the strategy was set and indications were being selected.

Biological risk came down by choosing a mechanical problem rather than a biological one. We were not guessing how biology would respond. We were building a machine, like any other machine, and machines yield.

What remained was engineering.

Why that residual and not another

One named risk, retained deliberately, chosen because it yields to the resource a founder can raise.

Money buys down engineering risk. It does not buy down biology. It does not create demand where none exists. It does not accelerate a clinical endpoint or shorten a regulatory pathway. Choosing engineering as the retained risk meant the thing we would spend years raising was the solvent for the thing that could kill us.

More precise than confidence. A claim about the match between the residual risk and the instrument available to reduce it.

The diagnostic value

Run the same inventory on the failures and it answers immediately, without hindsight.

A company carrying market risk, competitive risk, regulatory risk and financing risk at once has not made a bet. It has accumulated exposure. A company carrying substitute risk, regulatory risk and capital access risk simultaneously is in the same position with different labels.

The count and character of the residual risks is the diagnostic, and it is observable in advance. You do not need the outcome to run it. You need the construction.

What this changes about Elephant Risk

The Structural Literacy series has treated Elephant Risk as inherent and given: the product and market risk that defines entrepreneurship and cannot be systematized away. That is not quite right, and the modification matters. You choose which elephant you keep. The irreducible risk is not handed to a founder along with the idea. It gets selected, and selecting it well is the highest-leverage decision available at founding.

The same severity, turned inward

Six companies have been examined here without mercy. Sparing the seventh, because it is mine, would make the whole exercise worthless.

So: where Auris was fortunate.

The buyer appeared at the right moment. J&J's gap was foreseeable in shape and not in timing. Part IX says plainly that acquirer appetite belongs to the buyer rather than the target, and that applies here. A different quarter, a different internal program, a different competitive posture, and the same company with the same four answers meets a different market.

The engineering resolved. Gate 1B is the single irreducible risk in this framework. We retained it deliberately, and it came out. Retaining the right risk improves the odds and settles nothing. Biology could have defeated the mechanism late, and there would have been no framework to blame for it.

And one outcome proves nothing. A sample of one supports a narrower statement: the construction was deliberate, the reasoning is documented, and the risks were isolated in a stated order. Whether that ordering produced the result is the one question this paper says no framework can answer about anyone.

I would rather write that here than leave a reader to supply it.

The honest boundary

Deliberate isolation improves the odds. It guarantees nothing. A well-constructed venture can still lose to biology, to timing, or to a capital market that closes. Construction and outcome are separable, which is the point: construction can be assessed independently of how things turned out. That separability is what makes the approach teachable rather than a war story with a framework attached. One outcome can be luck. The same strategy, deliberately run a second time in an unrelated market, is method.

Part XI: Where This Sits in the Literature

This paper is intended to advance strategic thinking and address Porter, Christensen, and Moore as they apply, or do not, in an entrepreneurial setting. Each operates on a different gate, and running them in the wrong order, and failing to account for their intended audience and application is itself a way companies die.

Porter

So what does this framework claim that others leave unexplained?

Precedence, dispositiveness, and a time index. Porter names five forces and treats them as a simultaneous snapshot of an industry that already exists. No mechanism in the model says substitutes must be resolved before the other four are worth analyzing. No split between what can be researched and what must be built. And no way to handle a category that does not exist yet, which is where every company in Part VI stood at founding.

The framework does not claim to have found new forces. It claims the order in which they must be resolved, and that failing the first one ends the analysis.

There is also a problem Porter never claimed to solve. Five Forces is exquisitely sensitive to how you define the market, and offers no rule for choosing. Run it on "surgical robotics" and you get one answer. Run it on "peripheral lung nodule access" and you get the opposite. The anatomy-level unit of analysis does more work in this paper than any of the forces.

And a deeper limitation. Five Forces evaluates industry attractiveness, which is a capital allocation question. The right question if you run a diversified corporation deciding which businesses to be in. It is not the founder's question, because a founder does not select an industry.

Airlines are the textbook unattractive industry. Southwest did fine. Industry structure does not determine venture outcome, and Porter never said it did, though it is frequently taught as though he had.

Two consequences follow. The generic strategies assume a firm that can choose, and a new entrant cannot select cost leadership against an incumbent with manufacturing and service scale; offering it as an option is how a company's weak position gets written up as a strategy. And Five Forces prescribes position rather than sequence. Where to stand, not what to resolve first. Sequence is the whole content of this framework.

Even more than any industry dynamic, this paper claims the order of the evaluation.

Christensen

The closest fit and the most useful, because his own central distinction separates the two problem cases in this paper.

New-market disruption competes against non-consumption. It serves people who previously did without, because the alternative was too expensive, too invasive, or unavailable. Auris exactly, and the substitute test in his vocabulary: a patient with a peripheral nodule was not choosing between two products but between an inadequate biopsy and nothing.

Low-end disruption attacks with lower cost into segments the incumbent has overshot or does not want. SS Innovations exactly.

Here his framework earns its place, because it predicts the difficulty rather than describing it. Low-end disruption works only if an upmarket ramp exists: win the low end, then climb into the segments that matter. In medical devices, regulation turns that ramp into a wall. The incumbent cheerfully cedes markets it never served, so the beachhead generates no leverage on the market where the money is, and climbing requires a regulatory and clinical program the low-end revenue cannot fund.

A sharper statement of the financing circle in Part IX, and it comes out of his framework rather than this one.

And the case that matters most. Vicarious was not disruption at all. It was sustaining innovation: a better product, aimed at the incumbent's best customers, in the incumbent's core market. Christensen's central empirical finding is that entrants lose sustaining contests nearly always, because the incumbent has every incentive to fight and every resource to fight with.

The theory predicted that outcome in 1997.

One honest note. His predictive record has been contested, most prominently by Jill Lepore, and the disruption literature has been criticized for post-hoc case selection. He is cited here as a lens, not a law.

Moore

Crossing the Chasm operates entirely after Gate 1, which is its value and its danger.

The value is that it describes something the other two do not: the gap between early adopters buying a vision and the early majority buying risk reduction. Different customers, different purchase criteria, and a product that delights the first can stall completely at the second. In this field that gap is where a company with real clinical validation and twelve enthusiastic sites discovers that hospital number thirteen wants references, service guarantees and a proceduralist training pathway.

Two contributions land directly. **Whole product** is half of Gate 4 stated from the customer's side: the consumable is not only a capture mechanism, it is part of what makes the system buyable at all. And in this sector the chasm is per geography, which is why an installed base abroad does not constitute a crossing at home.

The danger sits in the assumption on his first page: that the product should exist. Apply Moore first and you get sophisticated go-to-market strategy for something nobody needs, which is what a failed Gate 1 company's investor deck looks like, and it is usually excellent work.

Sarasvathy

Effectuation is the literature closest to how the founding decisions in this paper were made, and the one least likely to be cited in a room full of medtech investors.

Her finding is that expert entrepreneurs do not begin with a market and search for a means to serve it. They begin with available means, who they are, what they know, whom they know, and work forward to what can be built, forming commitments with stakeholders as they go. The goal emerges from the process rather than preceding it.

A more accurate account of walking into a laboratory at Columbia, recognizing what flexible continuum robotics made possible, and building a company around the recognition, than any process model of opportunity identification.

It matters for a specific reason. This paper argues for rigorous prior analysis, and a reader could reasonably ask whether that is compatible with how founding happens. Sarasvathy's answer, and mine, is that the two operate on different questions. Effectuation describes how a founder arrives at a candidate. The gates describe how a candidate should be tested before capital is committed to it. Neither substitutes for the other, and a founder with only the first will build something impressive that nobody needs.

The synthesis

Each framework answers a different question, and the order is not optional.

Porter tells you whether the position is defensible.

Christensen tells you whether the attack vector is survivable.

Moore tells you how to sell it once you have one.

Run in that order they are complementary, and this paper adds a sequencing rule to a literature that mostly lacks one.

Run out of order they are how companies die well-advised: a beautiful go-to-market plan for an undefendable position, attacked along a vector the theory says entrants lose.

Part XII: What To Do About It

A diagnosis you cannot act on is entertainment. So this section assumes you have run the test and you did not like the answer.

First, work out which failure you have

Two companies can both fail the first gate today and require completely different responses.

You never passed. The substitute existed at founding and it exists now. Vicarious is this case: ventral hernia repair was routine in 2014 and it was routine in 2026. Nothing decayed, because there was nothing to decay.

You passed and the gate closed. The capability was unavailable when you started, and while you were building, someone else arrived. Every long-pathway device company is exposed to this, and it is no failure of judgment. The forecast horizon problem from Part II simply arrives on schedule.

The distinction matters because the second company has assets the first does not: a real technical achievement, a clinical claim that was true, and usually a portfolio that someone was designing around. Those are worth something. The first company's assets are worth what the components fetch.

Four moves

Reposition to unclaimed anatomy. Foreclosure in this field runs per anatomy, not across the board. A company that failed on abdominal soft tissue has not been told that surgical robotics is closed. It has been told that one target is taken. The question is whether the platform can be redirected to anatomy where the substitute test still passes, and the honest answer is usually that redirection costs more than founders expect and less than they fear. Auris was itself a platform bet across multiple applications, not a single-procedure company.

Narrow to a substitute-free pocket. Inside a served category there are usually specific presentations that nothing handles well: the difficult anatomy, the reoperative case, the patient the standard approach excludes. That is a smaller company than the one in the plan. It may also be the only real one available, and it is a legitimate answer to the first gate. The constraint is Gate 3, and Part VIII is the warning: a substitute-free pocket with insufficient volume is Medrobotics.

Sell as a defensive instrument. If an incumbent already owns the annuity in your category, your technology may be worth something to them as a way to stop losing it, independent of whether it would ever have built a company. Real transactions, and they happen regularly. Price it honestly as what it is. It is an Orthotaxy outcome, not an Auris outcome, and the difference is roughly two orders of magnitude. Founders who pursue this while telling their board they are pursuing the other one do everyone damage.

Stop. Occasionally correct, almost never on offer. A company with eighteen months of runway, no first-gate answer, and no repositioning path is going to spend that eighteen months and arrive at the same place with less. Returning capital is a legitimate act and the people who do it are disproportionately the people who get funded again.

What not to do

Do not raise against a failed first gate on the strength of the other three.

The single most expensive error in the record, and it does not look like an error at the time. The patent portfolio is real. The engineering is impressive. The market is large at the category level. Every one of those is true, and none of them touches the question of whether any substitute achieves the outcome you are proposing.

Vicarious raised for twelve years against exactly that combination.

The diagnostic

Five questions. If you cannot answer them in a sentence each, without reaching for slides, you do not yet have the answers.

1. Name the operation. What specifically becomes possible? Not better, not faster, not cheaper. Possible.

2. Name the substitute. What else achieves that same clinical outcome, at the same cost to the patient? If the honest answer is a list, price will be set by rivalry and you should plan accordingly.

3. Name the intellectual property. What lets you sell this without the incumbent's permission? And if the answer is a de novo pathway rather than a portfolio, say so, because that is a different and, perhaps, a stronger answer.

4. Name the number. How many systems are in demand, at what price, with what features? Not the category TAM. Units.

5. Name the consumable. What recurring high-margin revenue does each installed system produce, and what is that multiplied by the number in question four?

A board that asks these five questions annually will not prevent every failure. It will prevent the one where everybody is surprised.

For investors: fund the answer, not the company

A failed first gate is a decline. Nothing downstream rescues it, and across seven companies nothing ever did.

An unanswered first gate is not a decline. It is a smaller investment. Fund the answer and nothing else: a defined program to establish what operation becomes possible and whether anything else

achieves it, sized to that question, with development capital released only on the result. If the answer comes back that a substitute exists, the loss is the cost of the study rather than the cost of a company.

The distinction matters because the two states are indistinguishable in a pitch. A founder who cannot name the substitute may be facing a market that has none, or a market full of them. The first is the best investment in this paper. The second is Vicarious.

This is the same discipline the founder owes the company from the other side of the table. No development until the answer exists.

Sometimes there is no candidate to test, and the first capital funds the search for one. That is the same instruction at an earlier stage. It also gives the five questions above a use beyond annual review. Any one of them that cannot be answered in a sentence is a question worth funding by itself, and answering it is almost always cheaper than assuming it.

Founder's note. Every other claim in this paper rests on a document. One does not. It is my account of a company I founded, and it should be read as testimony rather than as evidence.

The first capital into Auris was raised to answer the first gate, and the question was broader than a single procedure. Not "is this operation unreachable," but "which operations are impossible today, by the human hand or with the instruments available, and where would enabling one produce a marked leap in outcome." That work was done over a defined period with outside parties assisting. No development proceeded until a procedure had been identified and the remaining gates asked of it.

The search took several candidates. The first was the inner ear. It was genuinely enabling, and it failed other gates described in this paper. The company is named for it. Auris is Latin for ear, and the name outlived the procedure. The second candidate was the back of the eye, and the company came close to being renamed accordingly before that one failed on gating problems of its own. The search continued until a procedure was found that cleared all four, and every procedure taken up afterward was stage-gated the same way.

Peripheral lung access was the output of that inquiry, not the premise of it. A flexible robotically driven instrument with no procedure attached is a technology without a market, and no amount of engineering fixes it. Until the operation was named, there was nothing to build toward.

This is not offered as proof that the staging rule works. One company is not a sample, and Part IX says so at length. It is offered because the rule is not hypothetical, and because a reader deciding whether to act on it is entitled to know that the author has.

Beyond surgery

Nothing in these four gates is surgical.

Substitutes, right to practice, demand, and capture are properties of any capital equipment business selling into an institutional buyer. The test travels, and the quickest way to see it travel is to run it on something adjacent.

Take domestic humanoid robots. Gate 1: what does the machine do that cannot be done another way? It folds laundry and cleans floors, both of which have well-established substitutes including the owner or a cleaning service. Gate 3: at eighty thousand dollars of capital equipment against fifty dollars of available household labor, the demand curve is elastic and the amortization period is long.

The framework returns a verdict in two sentences, and the verdict says nothing about the technology, which is remarkable. It is about substitutes and elasticity, which is where the money sits.

The value of running the gates in order is this. They tell you quickly which questions are worth spending years on.

Conclusion: you choose your elephant

Figure 6. What can be known before the money is spent, and what cannot.

The Structural Literacy series has kept two kinds of risk apart. Elephant Risk is the product and market risk that comes with the territory. Structural Risk is the governance, the capital structure and the deal mechanics that cost founders companies they should have kept.

This paper forces a modification to the first one, and the modification is the most useful thing in it.

Look at what the four gates are made of. Whether a substitute exists yields to research. How many units a market will take yields to research. Whether you can operate without permission yields to counsel and the filing record. Whether each installed system throws off recurring revenue is arithmetic. So Elephant Risk is not one thing. It is three, and only one of them is inherent.

Knowable before building. Existing substitutes, issued claims, procedure counts, reimbursement rules, incumbent portfolios, regulatory precedent. Public, dated, and checkable by anyone who bothers.

Boundable but not knowable. Future demand, attainable price, adoption speed, competitive response, coverage decisions not yet made, acquirer appetite. Research narrows these. Nothing settles them.

Discoverable only by building. Technical feasibility, workflow fit, whether the organization can execute.

The first category is not really Elephant Risk at all. It is preventable failure wearing the costume of inherent risk, and it behaves like Structural Risk in every way that matters: written down, checkable, and lost by not looking. The second yields to process in part and never entirely. Only the third is inherent, and it is the one worth keeping, because it is the one money can buy down.

The failures in this field are overwhelmingly failures of the first category. Not wrong estimates of the second, and not bad luck in the third. Nobody looked up what already existed.

So the irreducible risk in a venture does not arrive attached to the idea. You select it. The job is to isolate every risk that can be isolated and keep the one that yields to the resource you can raise, which in a machine business is engineering. Money buys engineering. It buys nothing else on the list. It does not buy down biology, it does not create demand, and it will not shorten a regulatory pathway or move a coverage decision by a single day.

You do not inherit your elephant. You pick it.

You do not inherit your structure either, and you pick it more than once.

The gates determine whether to build. They do not determine whether you keep what you build. That is a separate question on a separate clock. Gate selection happens at founding, and the record of it closes early. Structural Risk recurs: at every financing, every board seat, every set of protective provisions and liquidation preferences, and again at the exit, where terms written years earlier finally execute. A founder can answer all four gates correctly, build what nobody else could build, sell it for a number that reaches the trade press, and receive very little of it. Part 1 of this series, *Capturing Value at Exit*, documents one such case against the record.

And notice where Structural Risk sits in the taxonomy above. Knowable before. Charter documents, term sheets, board composition, the mechanics of an earnout: none of it is hidden, all of it is written down, all of it is checkable by anyone who bothers. The same category of failure this paper describes, on a shorter cycle, with more chances to make it.

Every company named in this paper was run by capable people. Most of the machines worked. In more than one case the engineering was better than the engineering that won.

They answered the questions out of order, and the order was sitting in public the whole time, on dates anyone could look up.

Go back to the lung. A nodule sits past the last branch a bronchoscope can follow, and somebody has to decide whether reaching it is a company or a paper. The four questions here will not tell you whether the machine can be built. They will tell you, before you spend a decade finding out, whether building it was ever the point.

Author Disclosure

Christopher Velis is the founder of Auris Surgical Robotics, Inc. (later Auris Health), which he incorporated in Delaware on 18 September 2007. He served as sole director, CEO, President, Treasurer and Secretary from incorporation through 2011. He licensed foundational technology from Columbia University and Johns Hopkins University, and recruited Frederic Moll as CEO in 2011.

He was not a shareholder of Auris Health at the time of its acquisition by Johnson & Johnson. He managed Grey Matter Health Ventures Fund I, LLC, a special purpose vehicle through which the company's earliest investors held former Auris Health securities, and received no compensation for managing it. Grey Matter's investors were among the stockholders represented by Fortis Advisors LLC in the litigation against Johnson & Johnson. The fund has since distributed the proceeds.

Auris appears in this paper as one of seven cases and is examined against the same framework as the others, including in Part X, where the elements of its outcome that the framework does not explain are set out directly.

Full documentation of founding provenance appears in Appendix A of *Capturing Value at Exit*, Part 1 of this series.

Disclaimer

This white paper is provided for educational purposes only and does not constitute legal, financial, or investment advice. The analysis rests on publicly available information and the author's interpretation of events. Founders, directors, and executives should consult qualified advisors before making decisions about capital commitments, regulatory strategy, or transaction structure.

Companies are named throughout; individuals are not. Where the paper characterises a company's position, it does so from filings and public record, and marks inference as inference. Where a claim rests on the author's direct experience, the text says so.

About AEIOU Academy

Why this exists.

Every other high-stakes profession has systematized its operational excellence.

Surgeons have protocols. Before an incision is made, there is a checklist, and preventable deaths have fallen by orders of magnitude. Pilots have checklists. Before every takeoff there is a run-up procedure, and flying has become the safest form of transportation on earth. Special operators have standard operating procedures, and after-action reviews examine every mission.

Entrepreneurship remains the last holdout. Founders learn through trial and error, mostly error. The knowledge that would prevent catastrophic mistakes exists, scattered across expensive professionals, hard-won experience, and lessons learned too late.

The approach.

AEIOU Academy provides the operating system for entrepreneurship: the systematized knowledge, frameworks, and practices that turn preventable failure into manageable risk.

The series distinguishes two kinds of risk. **Elephant Risk** is the irreducible risk in a venture, the part that can only be resolved by building. **Structural Risk** is the governance, capital structure, and deal mechanics that have nothing to do with whether the product works, and which can be engineered away through proper structure and practice.

This paper refines the first of those. Elephant Risk is not handed to a founder along with the idea. It is selected, and selecting it well is the highest-leverage decision available at founding.

AEIOU exists to eliminate Structural Risk, so that founder outcomes are determined by whether they built something valuable rather than by whether they understood the fine print.

About the Author

Christopher J. Velis, NACD.DC

Co-Founder, AEIOU Academy · Founder, Auris Health

Christopher Velis has occupied every seat at the table. He founded Auris Health, incorporating it in Delaware in 2007 and serving as sole director and chief executive through 2011, licensing the foundational technology from Columbia University and Johns Hopkins University before recruiting operational leadership to scale it. The company was acquired by Johnson & Johnson in 2019 for $3.4 billion at closing against a $5.75 billion headline, the largest acquisition of a venture-backed medical device company on record.

As an investment banker he advised on multiple billions of dollars in M&A transactions. As a venture investor he has evaluated hundreds of companies and founded a medtech and biotech fund. He holds the NACD Directorship Certification and is a Doctor of Business Administration candidate at Babson College, where his research examines decision-making.

Contact

Christopher J. Velis
velis@aeiouacademy.org
linkedin.com/in/chrisvelis
aeioumedical.org

Christos Kelepouris
christos@aeiouacademy.org
linkedin.com/in/christoskelepouris

Notes

Vicarious Surgical, capital and consumption. The figure used in the text is accumulated deficit of $253.4M as of 31 March 2026, the last balance sheet filed before dissolution, against $31.6M as of 31 December 2020. Accumulated deficit is unaudited at the interim date, is stated since inception, and measures cumulative losses rather than capital raised. Source: Form 10-Q for the quarter ended 31 March 2026, accession 0001213900-26-049749, and Form 424B3 filed October 2021. SEC EDGAR, CIK 1812173. A raised-to-date figure of $240.61M, including $3.50M of known debt across thirteen disclosed transactions, appeared in earlier drafts, cross-checked against PitchBook company profile 129143-71 retrieved 27 July 2026. It is not used here. That figure reconciles to disclosed rounds of $95.02M plus the $142M private placement plus $3.50M of debt, within approximately $90,000, but only on the assumption that the SPAC trust was fully redeemed. The merger Form 8-K shows otherwise: the trust delivered $77.59M. Adding that to the disclosed rounds, the private placement and the debt gives approximately $318.2M of total capital, of which only the $77.59M trust figure and the $142M private placement rest on primary filings. That hybrid is not used in the text because its largest component is a commercial compilation rather than a filing, and because it cannot be fully reconciled against the accumulated deficit and the $3.7M of cash remaining at the end. Business combination with D8 Holdings Corp. closed 17 September 2021 at approximately $1.0 billion in aggregate merger consideration. Announced 15 April 2021 at up to $460 million of gross proceeds assuming no redemptions, revised to approximately $487 million on 9 September 2021 after the private placement was upsized to $142M, and closed eight days later at $219.6M gross and $190.3M net of transaction costs of $29.3M. Redemption of approximately 78 percent of the trust is DERIVED from the $345M trust balance and the $77.59M delivered. Source: Form 8-K filed 23 September 2021, Exhibits 99.1 and 99.4, SEC EDGAR, CIK 1812173.

39 patent families, 188 documents, 63 active grants and 78 pending applications at dissolution. Forward citations by assignee: Intuitive Surgical 44, Globus Medical 37, Altec 31, Distalmotion 17, Cilag (Johnson & Johnson) 13. Patent analytics per PitchBook, derived from published EPO and USPTO records.

Stockholders approved dissolution and an assignment for the benefit of creditors on 21 July 2026; certificate of dissolution filed on or about 22 July 2026. The company reported no revenue in any fiscal year from incorporation in 2014.

Medrobotics Corporation, Chapter 7 petition filed January 2022, against asserted creditor claims exceeding $135M. The Flex Robotic System received FDA clearance in 2015 and held CE marking; contemporaneous trade reporting placed the installed base at approximately thirty hospitals globally.

CMS national coverage determinations on lumbar artificial disc replacement: CAG-00292N, decision memo 16 May 2006, issuing national noncoverage for LADR with the Charite disc for the Medicare population over sixty; and CAG-00292R, decision memo 14 August 2007, amending NCD Manual section 150.10 to noncoverage of the LADR procedure for that population. In the 2006 memo CMS stated it anticipated reopening the determination when another lumbar disc received FDA approval. ProDisc-L was approved 14 August 2006. Source: CMS Medicare Coverage Database.

Asensus Surgical; accumulated deficit of $987.6M as of 30 June 2024, the last balance sheet filed before the acquisition by Karl Storz. The merger agreement of 6 June 2024 provided $0.35 per share in cash; 272,616,330 shares were outstanding at 30 June 2024, deriving approximately $95.4M in common equity consideration. Accumulated deficit is unaudited and measures cumulative losses, not capital raised. Source: Form 10-Q for the quarter ended 30 June 2024, SEC EDGAR CIK 876378, accession 0001437749-24-026428.

Intuitive Surgical was founded in 1995 and did not obtain FDA clearance for the da Vinci system until 2000. Computer Motion held the earlier cleared robot for abdominal surgery. Intuitive reported a net loss of $18.4M in fiscal 2002 and $9.6M in fiscal 2003, and its first profitable full year was 2004, with net income of $23.5M: nine consecutive loss years from founding. Accumulated deficit stood at $128.8M at 31 December 2002. Total revenue for the year ended 31 December 2025 was $10.1 billion, an increase of 21 percent over $8.4 billion in 2024, with approximately 3,153,000 da Vinci procedures performed during the year and an installed base of 11,106 systems at year end. Sources: Intuitive Surgical Forms 8-K reporting fiscal 2003 and fiscal 2004 results; Form 10-K for the fiscal year ended 31 December 2025, filed 3 February 2026. All SEC EDGAR, CIK 1035267.

Section 513(i) of the Federal Food, Drug, and Cosmetic Act and 21 CFR 807.100. Note that a predicate device need not remain in commercial distribution, and a finding of substantial equivalence establishes a regulatory relationship between devices rather than clinical, economic, or competitive equivalence.

Section 513(f)(2) of the Federal Food, Drug, and Cosmetic Act. A de novo grant creates a new classification; subsequent devices may then use the 510(k) pathway by demonstrating substantial equivalence within that classification. It does not confer exclusivity.

SS Innovations International Inc., press release and Form 8-K, submission dated 5 December 2025, covering general, urological, colorectal, gynecological and cardiac indications. The company disclosed that the 510(k) route was selected following a pre-submission meeting. As of late July 2026 no clearance decision has been announced, past the company's own guidance of the first half of 2026.

Company disclosure of a cardiac telesurgery procedure conducted between Guyana and Indore, India, with reported round-trip latency of 290 to 300 milliseconds. Company-reported figures; not independently verified.

Earliest recorded filing 5 May 2014; family includes EP-3139843. The four next-most-cited filings in the portfolio are directed to a virtual reality surgical device, a virtual reality wrist assembly, and a virtual reality surgical camera system.

Intuitive Surgical, Inc. v. Computer Motion, Inc., 214 F. Supp. 2d 433 (D. Del. 2002); complaint filed April 2001. Computer Motion's earlier action asserting nine of its own patents was filed in the Central District of California in May 2000, with trial anticipated in 2003.

License agreement dated 22 December 1997, as described in the Delaware court's opinion at 214 F. Supp. 2d 433. Section 5.1 of the agreement grants Intuitive the right but not the obligation to enforce the LARS and ROBODOC patents within the field, at Intuitive's expense, with IBM's cooperation.

Twenty-year terms from filing on applications dating to the 1990s and early 2000s. INFERENCE: the timing of the 2020s entrant cohort is consistent with expiry of the foundational estate, but this is an inference from the pattern of entry rather than a sourced finding.

Johnson & Johnson acquisition of Auris Health, announced February 2019: $3.4 billion in cash at closing plus up to $2.35 billion contingent on ten regulatory and commercial milestones, a headline of $5.75 billion. None of the milestones was met. Fortis Advisors, representing the former stockholders, sued in the Delaware Court of Chancery; after a ten-day trial the court found breach of the efforts covenant and fraud and awarded over $1 billion. The Delaware Supreme Court affirmed in part and reversed on the implied covenant claim on 12 January 2026 (*Johnson & Johnson v. Fortis Advisors LLC*, 352 A.3d 229 (Del. 2026)), remanding for recalculation. On 26 January 2026 the Court of Chancery entered a final stipulated judgment of approximately $811 million, reported as the largest earnout award in Delaware history; Johnson & Johnson's Form 10-Q records payment of a

$0.8 billion judgment, inclusive of interest, in January 2026. Total consideration therefore reached roughly $4.2 billion against a $5.75 billion headline, with the earnout paying about 35 cents on the contracted dollar. Net proceeds to stockholders after legal fees and expenses are lower. The transaction is the subject of Part 1 of this series, *Capturing Value at Exit*.

FDA clearance for transoral robotic surgery for benign and selected malignant tumours of the mouth, pharynx and larynx, granted December 2009.

See note on Medrobotics above. Addressable procedure volume is now sourced; see the note on transoral robotic surgery volume.

VELYS Robotic-Assisted Solution, DePuy Synthes. FDA clearance January 2021 for total knee arthroplasty with the ATTUNE Total Knee System; unicompartmental indication 2024. Market and procedure figures are company-reported.

Orthotaxy SAS, acquired by Johnson & Johnson in 2018. Consideration not disclosed.

Company financial statements as filed. Total stockholders' equity across the same three periods: $46.7M, $9.8M, $3.6M. Cash provided by investing activities of $32.1M and $32.6M in the final two periods is consistent with liquidation of a securities portfolio to fund operations rather than an investment program.

Business combination with D8 Holdings Corp., completed September 2021. The Form 8-K check called for in earlier drafts has been completed and partly reverses the prior inference. Disclosed rounds of $95.02M plus the $142M private placement plus $3.50M of debt do reconcile to the $240.61M raised-to-date figure within approximately $90,000, but that reconciliation holds only if the trust delivered nothing. FACT: the trust delivered $77.59M, and gross proceeds were $219.6M. The trust was heavily redeemed at approximately 78 percent, not substantially fully redeemed, and net proceeds exceeded the private placement alone by roughly $48M.

Becton, Dickinson and Company. Recorded as a new investor in the September 2021 placement, a seller in the August 2023 secondary offering, and as holding no position at dissolution.

The da Vinci SP system received an initial clearance in April 2014 and a further clearance on 31 May 2018 for urologic procedures, with US launch in the third quarter of 2018. Clearances for inguinal hernia repair, cholecystectomy and appendectomy followed in December 2025. The single-port platform did not hold the indication relevant to this case until well after the period in question, which is why the argument here rests on Gate 1(a) at founding rather than on this date.

FDA premarket approval records. Charité (DePuy Spine) approved 26 October 2004 for single-level replacement; ProDisc-L (Synthes Spine) approved 14 August 2006.

Centers for Medicare & Medicaid Services national coverage analysis, tracking sheet CAG-00292R. Determination issued 16 May 2006; reconsideration opened November 2006; decision memorandum issued 14 August 2007.

Charité withdrawn from commercial distribution in 2012. activL approved June 2015. For context, cervical devices followed a different path: Prestige approved 16 July 2007 and ProDisc-C approved 17 December 2007.

Published phase-transition success rates by therapeutic area are widely used in pharmaceutical valuation. NO SPECIFIC SOURCE CITED: if this comparison is retained, a standard reference such as the BIO/Informa clinical development success rates study should be attached.

Michael E. Porter, Competitive Strategy: Techniques for Analyzing Industries and Competitors (Free Press, 1980); and 'How Competitive Forces Shape Strategy', Harvard Business Review, March 1979.

Clayton M. Christensen, The Innovator's Dilemma (Harvard Business School Press, 1997); and Christensen and Raynor, The Innovator's Solution (2003), for the distinction between new-market and low-end disruption.

Jill Lepore, 'The Disruption Machine', The New Yorker, 23 June 2014.

Geoffrey A. Moore, Crossing the Chasm (HarperBusiness, 1991).

Saras D. Sarasvathy, 'Causation and Effectuation: Toward a Theoretical Shift from Economic Inevitability to Entrepreneurial Contingency', Academy of Management Review 26(2), 2001; and Effectuation: Elements of Entrepreneurial Expertise (Edward Elgar, 2008).

Illustrative. Price point and substitute labor cost are stated for argument and are not drawn from a specific product or market study.

INFERENCE, not a sourced finding, and the most contestable claim in this paper. It follows from installed-base arithmetic: procedure volume generates training data, and the incumbent's volume compounds annually against an entrant's. The counterargument is that model architecture, synthetic data, or externally sourced datasets could decouple performance from proprietary volume. That has not happened in this field to date. Read the claim as a prediction.

National Cancer Database analysis of transoral robotic surgery adoption, 2010 to 2016: of 73,661 patients with oropharyngeal squamous cell carcinoma, 50,643 were treated nonsurgically, 18,024 with non-robotic surgery and 4,994 with TORS; annual TORS cases rose from 363 in 2010 to 994 in 2016. The study defines a high-volume center as ten or more cases per year. Note that these figures describe the addressable procedure pool rather than Medrobotics' own installed base, and that the Flex system competed for that pool against an incumbent already cleared for the indication.

Drawn from the company's own announcements rather than from the submission itself, which is not public. SS Innovations International reported a 5 December 2025 filing covering general surgery, urology, colorectal, gynecology and cardiac indications for the SSi Mantra system, and stated separately that telesurgery authorization would be pursued following FDA and CE marking of the system. The scope of the filing as submitted has not been independently verified. Whether long-distance telesurgery could be reached through a substantial-equivalence pathway at all is a separate and open question.

FDA 510(k) K203614, Auris Health, cleared March 2021 for the reprocessed Monarch Bronchoscope, described in the summary as a reprocessed single-use device validated for one reprocessing cycle. The GUDID record for the Monarch Bronchoscope lists the device as single use and supplied sterile. The scope carries a working channel and four-way articulation and is used with the Monarch Platform under K193534.

Intuitive Surgical quarterly and annual results as filed. FY2001 revenue $51.7M and FY2002 revenue $72.0M, both at a net loss; Q2 2003 first profitable quarter with net income of $0.9M against a $3.7M loss in Q2 2002; FY2003 revenue $91.7M with a net loss of $9.6M; FY2004 revenue $138.8M, up 51 percent, with net income of $23.5M. Recurring revenue rose from $7.0M in 2001 to $60.0M in 2004, and from 29 percent of sales in Q1 2003 to 46 percent in Q1 2004. Cumulative systems sold: 149 at end 2002, 210 at end 2003, 286 at end 2004. The Computer Motion acquisition closed 30 June 2003. The attribution of the turn to prostatectomy adoption is INFERENCE from the coincidence of timing with the shift in recurring revenue and utilization, and from contemporaneous accounts; it is not established by the filings themselves.

Yaxley JW, Coughlin GD, Chambers SK, et al. Robot-assisted laparoscopic prostatectomy versus open radical retropubic prostatectomy: early outcomes from a randomised controlled phase 3 study.

Lancet 2016;388(10049):1057-1066. 326 men enrolled between August 2010 and November 2014, randomised 163 to each arm; no significant difference in urinary or sexual function at six or twelve weeks. Twenty-four month outcomes reported in Coughlin GD, Yaxley JW, Chambers SK, et al. Lancet Oncology 2018;19(8):1051-1060, again finding similar functional outcomes. The adoption figures are from the accompanying Lancet comment, 'Robotic surgery evaluation: 10 years too late', which records the rise from approximately 1.8 percent to 85 percent of US radical prostatectomies between 2003 and 2013 in the absence of high-level comparative evidence.

Hansen Medical Inc., Form 10-Q filings. IPO completed 15 November 2006, 7,187,500 shares, net proceeds $78.3 million. CE marking for the Sensei system in the fourth quarter of 2006; FDA clearance for Sensei and Artisan catheters and first product revenue in the second quarter of 2007. Subsequent offerings: 3,000,000 shares for approximately $39.5 million net on 7 April 2008; 11,692,000 shares for approximately $35.3 million on 22 April 2009; 16,100,000 shares for approximately $29.8 million on 20 April 2010. Acquired by Auris Surgical Robotics, announced 19 April 2016, at $4.00 per share in cash, total equity value approximately $80 million, a premium of roughly 39.9 percent to the prior closing price.