

Lock-In Cost

Arc 2 · vocabulary · Portfolio company

In the room: Operating partner, portfolio CEO, anyone approving an automation spend
The question: What is the deferral actually costing, and why does the call never get made?

THE POSITION

SIGHTING

An automation programme is approved. Eighteen months later everyone agrees the design was wrong, and nobody moves, because the correct alternative now costs three times what it would have before the build.

THE CONSENSUS READ

Sunk cost. Ignore what's spent, evaluate forward, make the right call now.

THE MECHANISM

Sunk cost is the wrong frame and it is why the call never gets made. What's spent is gone and irrelevant. What matters is that the price of the alternative rose — integrations, retrained staff, a data model everything depends on, a vendor contract with term left. That is not sunk. It is a live and climbing liability.

THE EXPOSURE

A fund/kill decision is an option that decays, and by month twenty-four it can go negative.

THE TEST

On one live automation, price the alternative today and price what it would have cost at approval. The gap is the number.

Lock-in cost

What the correct alternative costs more than it would have cost before you committed.

The test. Price the second-best topology today, and price it as of the approval date. If the gap is inside normal estimating error, you have no lock-in and this doesn't apply. If it's a multiple, the decision you're deferring is getting more expensive on a clock you're not watching.

THE BRIDGE — CLAIMED TO REALISED

Sizing is applied to the same composite programme as Episode 01. The Staw and Dietvorst findings are real, published and cited with their samples; the dry-powder figure is internal market research and flagged as such.

LINE	MOVES	RUNNING
Claimed — A fund/kill decision deferred at month 18 to 'gather more data', at no cost		\$0 to defer
1. Deferral is priced at zero, and that is the error [changes the story]	-\$0	decision valued at \$0 to defer
The decision to wait is almost never written down as a decision.		
2. What actually rises	-\$280,000 to -\$450,000	-\$280,000 to -\$450,000
Four components, all estimable in an afternoon.		
3. The decision quality degrades on the same clock	-\$140,000 to -\$190,000	-\$420,000 to -\$640,000
The half nobody prices, and it has a measured basis.		
4. Where I was wrong — I priced lock-in as technical debt [where I was wrong]	-\$310,000 to -\$420,000	-\$730,000 to -\$1.06M
My first model of lock-in cost was the first two components of Line 2: integration surface plus rebuild. Cost: Roughly half the true figure, understated in a board paper.		
5. The option curve [residual]	+\$0	-\$730,000 to -\$1.06M
Put the three together and the fund/kill option has a shape worth drawing on one slide.		
Realised	A decision worth its full value at month 0 is worth \$730,000 to \$1.06M less by month 24	

THE DOWNSIDE

Multiple: 9x, illustrative

\$730,000–\$1.06M of accumulated lock-in and trust erosion at an illustrative 9x, against a decision that cost nothing to make at month zero and was never framed as a decision at all. Substitute your own comps.

\$6.6M to \$9.5M of enterprise value — against a decision nobody recorded making

When it surfaces. Never, clearly — which is what makes this the hardest of the ten to act on. Lock-in cost doesn't appear as an event. It appears as a programme everyone privately agrees was a mistake and nobody proposes stopping, booked as ongoing operating cost until an exit process forces the question.

Lock-In Cost

The question: What is the deferral actually costing, and why does the call never get made?
Paste into a request list or a management agenda. Each question resolves to an artifact, not to a characterisation.

THE DILIGENCE PACK

1. For each automation programme over \$250K, provide a current estimate to implement the second-best structural alternative, and a reconstruction of what that estimate would have been at approval.

Artifact: Two estimates, same scope, same team.
The gap is lock-in cost and nobody has ever asked for it.

2. List every downstream system, report and KPI whose definition depends on how the automated system structures its records.

Artifact: A dependency list from data/BI, not from the programme team.
This is the component that surprises people.

3. For any programme where a review has been deferred, provide the written record of the deferral, who requested it, and what additional data was commissioned.

Artifact: The meeting records.
Per Staw, post-threshold information-gathering is the escalation response. Read it as a diagnostic, not as diligence.

4. Identify who approved the original programme and who will present the fund/kill review.

Artifact: Names.
If they're the same person, the review has a known and measured bias.

5. State the date and metric at which this programme stops, and name the person who will read that metric — who may not be the approver.

Artifact: One line in the approval memo.
Free, and the only control here that works.

DISQUALIFIER

If nobody can produce estimate 1, the programme has never been compared to an alternative — which means the original approval was a yes/no on one option rather than a choice between two.

SOURCES — EVERY MEASUREMENT WITH ITS SAMPLE

CLAIM	SOURCE	SAMPLE / CLASS
~75% of subjects responsible for a prior failure sought retrospective justifying information, against ~25% of those not responsible	Staw, escalating commitment literature (1976 onward)	Experimental subjects; responsibility manipulated (measured)
Participants who watched a model err were significantly less likely to bet on it, including the 83% who watched it beat the human; seeing a human err did not reduce willingness to use the human	Dietvorst, Simmons & Massey, Journal of Experimental Psychology: General 144(1), 2015	Five studies, incentivised choice (measured)
Lock-in accumulation of \$280,000–\$450,000 over 18 months; trust erosion of \$310,000–\$420,000 annualised	Applied to the same composite programme as Episode 01	n/a — composite, our sizing (composite)
Roughly \$1.3 trillion of global buyout dry powder outstanding, concentrated in aging vintages under pressure to deploy	Internal market research	n/a — not a published measurement; held loosely, direction only (reported)
9x entry multiple	Illustrative mid-market comp	n/a — illustrative, arithmetic exposed (illustrative)

One standing caveat. Every number in this show is somebody else's measurement, and I'll tell you whose, with the sample. None of it is diligence on your deal. Do that yourself.