19 *The Economics of Automation and AI*

After this chapter you will be able to:

1. recognize an automation or AI proposal as an ordinary challenger, and compare it with the incumbent process by annual equivalent cost
2. build the installed cost basis from the vendor quotation, the integration cost, and the retraining cost, leaving sunk studies out
3. convert each claimed benefit into the cash the firm actually stops paying, and write an explicit zero when a benefit fails that test
4. diagnose a vendor ROI claim at both ends of the fraction: the bare-quote denominator and the full-utilization numerator
5. decide between adopting now and waiting one year with the rule $(AEC_1 - AEC_2)/i > AEC_D - AEC_1$
6. compute the unit cost of the n-th deployment from the learning curve $C_n = C_1\, n^{\log_2 \varphi}$

Every generation of engineers meets a machine that promises to do the work of people, faster and around the clock, and every generation asks the same question: is it worth the money? The current names are industrial robots, autonomous haulers, and AI systems that watch, predict, and decide. The names will change again. The economics will not. An automation proposal is a *challenger* in the exact sense of chapter 16: a new way to deliver a service the firm already delivers, priced against the way the firm delivers it now. This chapter applies the book's machinery, the annual equivalent cost of chapter 6, the defender–challenger comparison and marginal rule of chapter 16, and the sensitivity discipline of chapter 12, to that technology class. What is different about automation gets its own treatment: a cost basis dominated by integration rather than hardware, benefits that arrive as avoided costs rather than revenue, vendors who quote returns nobody should accept unaudited, challengers that improve while you deliberate, and costs that fall with experience along a learning curve. The running case is an autonomous inspection robot at *Najd Copper*, the concentrate producer this book has followed since Chapter 5.

What a payback claim leaves out.

Q1. A vendor quotes 220,000 SAR for an inspection robot and claims it "pays for itself in six months." Name two numbers you would need to see before believing either the cost or the saving.

Q2. Najd Copper currently inspects its conveyor gallery with two technicians. A robot could take over the task. In the vocabulary of chapter 16, what is the defender, what is the challenger, and what is the defender's "market value"?

Q3. The first robot deployment costs 10,000 SAR to commission, and the learning rate is 90%. What does commissioning cost for the second unit? For the fourth?

Answers.

Q1. On the cost side: the integration and retraining spending needed to make the robot useful on *this* site, which the quote does not include. On the saving side: the utilization assumption behind the claimed saving (what fraction of the planned duty the robot will actually achieve), and which of the claimed savings are cash the firm currently spends. A payback claim built on the bare quote and a 100% duty cycle fails on both ends.

Q2. The manual inspection process is the defender; the robot is the challenger. The defender's market value is essentially zero: there is nothing to sell, so keeping the current process forgoes no cash. Its annual equivalent cost is simply what the firm spends each year to run it. The comparison is the standard one: least annual equivalent cost wins.

Q3. Each *doubling* of the count multiplies the unit cost by the learning rate: the second unit costs $10{,}000 \times 0.90 = 9{,}000$ SAR, and the fourth costs $10{,}000 \times 0.90^2 = 8{,}100$ SAR. The clock of a learning curve is the doubling, not the calendar year.

Reference. The replacement machinery reused throughout this chapter is developed in chapter 16; Park, 3rd ed., Chapter 12.

Automation challenger. A machine, robot, or software system proposed to deliver a service the firm already delivers by other means. It is evaluated exactly like any challenger: by annual equivalent cost against the incumbent process.

19.1 *Automation Is a Challenger*

Strip away the vocabulary of the sales brochure and an automation proposal has a familiar shape. The firm already gets the work done: technicians walk the gallery, loaders stack the bags, planners schedule the maintenance. That incumbent process is the *defender*. The proposed robot or AI system is the *challenger*. Both deliver the same required service, so they are *service projects* in the sense of Chapter 5: compare annual equivalent costs and take the cheaper. Nothing in this chapter suspends a single rule from the earlier chapters; "should a machine replace labor, and when?" is the oldest question in engineering economy, and AI is only its current face.

Two features of the incumbent process simplify the comparison. First, a manual process usually has no resale value, so the outsider viewpoint of Section 16.2 prices the defender at zero investment: its annual equivalent cost is just its yearly spending. Second, that yearly spending is exactly the list of things the challenger promises to remove, so the defender's cost lines and the challenger's benefit lines are the same numbers seen from opposite sides. Keeping the two views consistent is half the discipline of this chapter; the other half is resisting numbers that appear on only one side of the ledger.

> **Insight.** Treating automation as "new economics" is how bad proposals get approved. Once the robot is recognized as an ordinary challenger, every protection the book has built applies automatically: sunk pilots do not count, the comparison is by annual equivalent cost, benefits must be cash flows the firm actually stops paying or starts receiving, and the verdict is re-run as prices and estimates move. The technology is new; the arithmetic is Chapter 16's.

What, then, deserves a chapter? Five practical features that automation challengers share and ordinary machine challengers mostly do not: the true cost basis is dominated by integration; the benefits arrive as engineering metrics that must be converted to cash honestly; the vendor supplies a ready-made analysis that is systematically flattering; the challenger itself improves and cheapens while the firm deliberates; and unit costs fall with each deployment along a learning curve. The sections that follow take these in order, then assemble the whole toolkit on the inspection robot.

Installed cost basis. The full investment required to make an automation challenger deliver its service: purchase price plus integration plus retraining. This is the I in every capital recovery computation.

Integration cost. The site-specific spending, communications, fixtures, safety systems, software connections, commissioning, needed to make purchased hardware work in one particular plant.

In practice. Plan the site survey before trusting any budget. Integration spending is set by the plant, not the product: wireless coverage in a steel gallery, the state of the equipment registry the new system must talk to, and how far the safety case has to be rebuilt. On brownfield sites the integration line routinely exceeds the license or hardware line, and only an engineering walk-down can price it.

19.2 *The Challenger's Real Cost Basis*

The price on the vendor's quotation is where the cost basis starts, not where it ends. A robot in a crate delivers no inspections. Before it earns a single riyal of savings, the plant must give it a map and a network, docking stations and safety interlocks, connections to the maintenance system, and people trained to run and repair it. All of that is investment made to obtain the service, so all of it belongs in the cost basis I that the capital recovery charge of chapter 6 spreads over the asset's life. The same principle set the cost basis for depreciation in Section 10.1: the basis is the cost of putting the asset *into service*, not the cost of the box.

Example 19.1 — What the inspection robot really costs

Najd Copper is offered an autonomous inspection robot for its conveyor gallery and confined plant spaces. The vendor's quotation reads 220,000 SAR for the robot, sensors, and a five-year software license. The project engineer's walk-down adds the site work: gallery wireless upgrades, docking stations, safety interlocks, and integration with the maintenance system, 140,000 SAR in total; and retraining, four technicians at 6,000 SAR each plus 16,000 SAR of backfill while they train, 40,000 SAR in total. Last year the firm also paid 25,000 SAR for a feasibility study of gallery automation. Determine the installed cost basis.

Solution.

Take the amounts one at a time, with the tests of chapter 16 in hand.

item	amount (SAR)	verdict
vendor quotation	220,000	**yes** — price of the hardware and license
integration (site work, comms, interlocks)	140,000	**yes** — needed to obtain the service
retraining (fees plus backfill)	40,000	**yes** — needed to obtain the service
feasibility study, last year	25,000	**no** — sunk
installed cost basis I	**400,000**	

Table 19.1. The inspection robot's installed cost basis. The vendor's quotation is barely more than half of it, and the sunk study is excluded entirely.

The feasibility study is spent whatever the firm decides; it is the 85,000-SAR truck of Example 16.1 again, and it stays out. The integration and retraining lines are future cash required to obtain the service, so they are investment:

$$I = 220{,}000 + 140{,}000 + 40{,}000 = \boxed{400{,}000 \text{ SAR}}.$$

The site work and retraining, 180,000 SAR together, come to about 82% of the vendor's quotation. Every capital recovery charge in the rest of this chapter is computed on the 400,000, not the 220,000, and Example 19.3 shows what happens to a payback claim when this distinction is ignored.

19.3 *Converting Benefits Honestly to Cash Flows*

Automation benefits arrive from the vendor as engineering metrics: cycle time down, scrap down, no injuries, runs all night. None of these is a cash flow yet. The conversion rule is the incremental principle of Chapter 1: a benefit is worth exactly the cash the firm *stops paying or starts receiving* because of it, in this plant, at this demand. The four classic metrics convert as follows.

- **Cycle time.** Faster cycles create capacity. Capacity is worth cash only if it displaces spending the firm currently makes (overtime, weekend shifts, purchased outside work) or produces output the market will actually buy. Unsold capacity is worth zero.
- **Scrap and quality.** A scrap reduction is worth the full cost of the scrapped units avoided, at the *measured* improvement, not the brochure improvement.
- **Safety.** Removing people from a hazard saves the measurable spending attached to the exposure: permits, standby crews, insurance premiums, claims history. The human value of an avoided injury is real and larger, but the cash line should carry only what the ledger supports, with the rest argued in words.
- **Availability (24/7).** A machine that can work nights is worth cash only if the nights are needed: sold demand, or a bottleneck relieved. An unmanned shift with nothing to make is a brochure line, not a cash flow.

Honest conversion. Pricing each claimed benefit at the cash the firm actually stops paying or starts receiving, at measured performance and current demand, and writing an explicit zero when a benefit fails the test.

Example 19.2 — Four metrics, four cash lines

Wadi Phosphate's workshop is offered a robotic deburring cell, installed cost 190,000 SAR, salvage 20,000 SAR after 5 years, operating cost 12,000 SAR per year, at $i = 10\%$. The vendor's letter claims 384,000 SAR per year of benefits. The workshop's own data: demand is 45,000 parts per year; the manual cycle is 6 minutes and the robot's is 4; the station's two shifts provide 4,000 hours per year, so at 6 minutes per part (4,500 hours) the shop chronically runs about 500 hours per year of overtime at a 90-SAR premium; measured scrap in the pilot fell from 3.0% to 2.0%, and a scrapped part costs 60 SAR; the deburring bench's injury and PPE history costs about 8,000 SAR per year, and safety engineering expects half of it to remain; there is no demand for a third shift. Convert each metric to an honest cash line and screen the cell.

Solution.

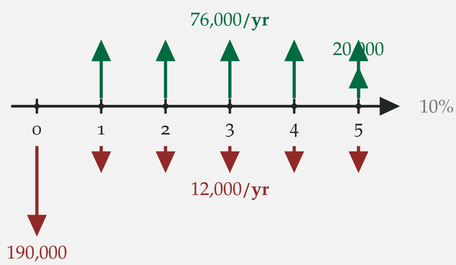

Figure 19.1. The deburring cell after the honest conversion: installed cost, five years of net honest benefits against operating cost, and a salvage recovery.

Step 1: convert metric by metric.

The conversion table places the vendor's line and the honest line side by side.

metric	honest mechanism	vendor (SAR/yr)	honest (SAR/yr)
cycle time 6 → 4 min	displaces 500 h/yr of overtime at 90 SAR	150,000	45,000
scrap 3.0% → 2.0%	45,000 × 0.010 = 450 parts at 60 SAR	54,000	27,000
safety	measured history 8,000, half remains	120,000	4,000
24/7 availability	no third-shift demand exists	60,000	**0**
total		384,000	**76,000**

Table 19.2. Converting the deburring cell's four claimed metrics to cash. The vendor priced

created capacity at full margin, took the brochure scrap delta, monetized injuries generously, and sold a night shift with no demand. The honest column keeps only cash the firm stops paying.

The vendor's cycle-time line prices 22,500 extra parts nobody has ordered; the honest line keeps only the overtime the speed-up actually displaces. The measured scrap delta is half the brochure's. The safety line keeps the ledger history. The explicit zero on the last row is the most important entry in the table.

Step 2: screen at the honest benefits.
With $(A/P, 10\%, 5) = 0.26380$:

$$CR = (190{,}000 - 20{,}000)\,(0.26380) + 0.10\,(20{,}000) = 44{,}846 + 2{,}000 = 46{,}846,$$
$$AEC = 46{,}846 + 12{,}000 = 58{,}846 \text{ SAR per year.}$$

Against honest benefits of 76,000 SAR per year, the net annual advantage is

$$76{,}000 - 58{,}846 = \boxed{+17{,}154 \text{ SAR per year}} \;\Rightarrow\; \text{accept.}$$

The cell survives the honest conversion, with a modest margin. At the vendor's 384,000 SAR the project looked like a five-month miracle; the honest case is a sound, ordinary machine purchase, and that is the version worth signing.

Insight. Write the zeros. An honest conversion table gains its authority from the rows it refuses to monetize: the unsold capacity, the unmanned shift, the un-measured quality delta. A reviewer who sees explicit zeros knows the other rows were tested the same way; a reviewer who sees every metric monetized knows none of them were.

Vendor ROI. A return or payback figure computed by the seller of the equipment. It is a sales document with the arithmetic of an analysis, and it is audited, never adopted.

19.4 *The Vendor-ROI Problem*

Automation proposals rarely arrive naked; they arrive wearing a vendor's return-on-investment sheet. "ROI in six months" is this technology class's signature claim, and it is manufactured the same way every time: the denominator is the bare quotation (no integration, no retraining), and the numerator assumes

100% utilization, every claimed metric monetized at brochure values, and no discounting. Each choice is individually defensible-sounding and collectively decisive. The audit procedure is the one this book has practiced since Chapter 5: rebuild the cost basis, re-convert the benefits honestly, restate at the MARR, and only then read the verdict.

In practice. Utilization has an engineering ceiling before it has an economic estimate. A robot that needs four hours of charging and a weekly maintenance window cannot exceed about 85% duty; a cell inside a line inherits the line's takt time and cannot run faster than the process feeding it. When a vendor's sheet assumes 100% utilization, the audit does not need to reach the economics to begin rejecting it.

Example 19.3 — Auditing "payback in six months"

The inspection-robot vendor's proposal to Najd Copper closes with this paragraph:

> *"Your investment: 220,000 SAR. The robot patrols 24/7, replacing the equivalent of five inspection technicians (5 × 96,000 = 480,000 SAR per year). Payback:* $220{,}000/480{,}000 \times 12 = 5.5$ *months. The system pays for itself twice in the first year."*

Audit the claim using the installed basis of Example 19.1 and the plant's own estimates: two technicians (48,000 SAR each) actually assigned to inspection; three unplanned stoppages per year at 40,000 SAR each, of which the reliability engineer judges half detectable early by nightly patrols; confined-space entry overheads (permits, standby crew, insurance) of 18,000 SAR per year; robot operating cost 30,000 SAR per year; engineering duty ceiling 85%; five-year life, salvage 40,000 SAR, $i = 10\%$, $(P/A, 10\%, 5) = 3.7908$.

Solution.

Flaw 1: the denominator is not the investment.

The 220,000 SAR quotation omits 180,000 SAR of integration and retraining. The installed basis is 400,000 SAR (Example 19.1), nearly double the denominator of the payback fraction.

Flaw 2: the numerator monetizes people who do not exist.

"The equivalent of five technicians" prices hypothetical night shifts the plant never staffed. The cash the firm can actually stop paying is two technicians, $2 \times 48{,}000 = 96{,}000$ SAR per year, and only under the labor conditions of Section 19.7.

Flaw 3: the honest benefits are smaller and must be built line by line.

Nightly patrols catch the detectable half of the stoppages, $0.5 \times 3 \times 40{,}000 = 60{,}000$ SAR per year; the confined-space overhead adds 18,000. The full-duty benefit pool is

$$96{,}000 + 60{,}000 + 18{,}000 = 174{,}000 \text{ SAR per year,}$$

and the duty ceiling scales it: at 85%, $174{,}000 \times 0.85 = 147{,}900$ SAR per year.

Flaw 4: no discounting.

Restate at the MARR. With $(A/P, 10\%, 5) = 0.26380$:

$$CR = (400{,}000 - 40{,}000)(0.26380) + 0.10\,(40{,}000) = 94{,}968 + 4{,}000 = 98{,}968,$$
$$AEC = 98{,}968 + 30{,}000 = 128{,}968 \text{ SAR per year.}$$

Net advantage at the 85% ceiling: $147{,}900 - 128{,}968 = 18{,}932$ SAR per year, worth $18{,}932 \times 3.7908 \approx 71{,}767$ SAR of present worth over the five years. The honest payback is $400{,}000/147{,}900 = 2.7$ years, not 5.5 months.

Verdict.

Reject the memo; accept the project. The robot is a sound investment on honest numbers, positive present worth at the MARR, but it is a 2.7-year machine, not a 6-month miracle. A firm that budgets on the vendor's sheet will approve the purchase and then starve the 180,000 SAR of integration that makes it work. The proposal deserved funding; the paragraph deserved the audit.

Insight. Vendor ROI fails at both ends of the fraction at once: the denominator omits integration and retraining (too small), and the numerator monetizes brochure metrics at full utilization (too large). The two errors multiply, which is how 2.7 years becomes 5.5 months. The repair is never to haggle over the vendor's figure; it is to rebuild the fraction from the installed basis and the honest conversion table.

19.5 *Fast Obsolescence: Buy Now or Wait?*

Ordinary challengers hold still while the firm deliberates; automation challengers do not. Next year's robot will be cheaper and better, and everyone in the approval meeting knows it. The temptation is to wait forever; the opposite temptation is

Challenger obsolescence. The improvement of the best available challenger over time: falling prices, better performance, shorter economic lives. It makes "adopt now or wait" a comparison between today's challenger and next year's, not only between defender and challenger.

to ignore the improvement and apply the marginal rule of Section 16.5 as if the challenger were frozen. Both are errors, and the arithmetic that separates them is short.

Let AEC_D be the defender's annual cost (the manual process), AEC_1 the best challenger's annual equivalent cost today, and AEC_2 the improved challenger's annual equivalent cost if the firm waits one year. Under repeatability (Section 16.7), adopting now commits the firm to a cost stream worth AEC_1/i as a capitalized equivalent (Section 5.7); waiting costs one more defender year, then AEC_2/i. Waiting is worthwhile exactly when

$$\underbrace{\frac{AEC_1 - AEC_2}{i}}_{\text{capitalized value of the improvement}} > \underbrace{AEC_D - AEC_1}_{\text{one year of forgone savings}}, \qquad (19.1)$$

that is, when the improvement, capitalized, outweighs the year of savings given up. Both sides are future flows; nothing sunk enters.

Example 19.4 — One robot generation is worth waiting for

Najd Copper's manual inspection process costs $AEC_D = 174{,}000$ SAR per year (Example 19.3, taken here at full duty for clarity). Today's robot has $AEC_1 = 128{,}968$ SAR per year. The vendor's roadmap prices next year's model at 340,000 SAR installed (the hardware cheapens; the integration does not), salvage 34,000 SAR, same 30,000 SAR operating cost, five-year life, $i = 10\%$. Should the firm adopt now or wait one year?

Solution.

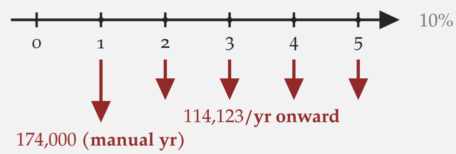

Figure 19.2. The "wait one year" option: one more year of the manual defender, then the improved challenger's annual cost stream from year 2 onward (continuing beyond the diagram under repeatability).

Step 1: the improved challenger's AEC.

With $(A/P, 10\%, 5) = 0.26380$:

$$CR_2 = (340{,}000 - 34{,}000)(0.26380) + 0.10\,(34{,}000) = 80{,}722.80 + 3{,}400 = 84{,}122.80,$$
$$AEC_2 = 84{,}122.80 + 30{,}000 = 114{,}122.80 \text{ SAR per year.}$$

Step 2: apply the wait rule.

The improvement is $AEC_1 - AEC_2 = 128{,}968 - 114{,}122.80 = 14{,}845.20$ SAR per year; the forgone saving is $AEC_D - AEC_1 = 174{,}000 - 128{,}968 = 45{,}032$ SAR. Test Equation (19.1):

$$\frac{14{,}845.20}{0.10} = 148{,}452 > 45{,}032 \Rightarrow \boxed{\text{wait one year}}.$$

Step 3: confirm by capitalized equivalents.

Adopt now: $CE = 128{,}968/0.10 = 1{,}289{,}680$ SAR. Wait: one defender year plus the improved stream, discounted one year at $(P/F, 10\%, 1) = 0.9091$:

$$CE = \big(174{,}000 + 114{,}122.80/0.10\big)(0.9091) = 1{,}315{,}228\,(0.9091) \approx 1{,}195{,}674\,\text{SAR}.$$

Waiting is cheaper by about 94,000 SAR of present worth. Please note what the naive rule would have said: $AEC_D = 174{,}000 > AEC_1 = 128{,}968$, so Section 16.5 taken alone shouts "replace now." The improving challenger changes the verdict, and Equation (19.1) says by how much.

Insight. The wait rule has a built-in stop against waiting forever. Each year of waiting costs the forgone saving $AEC_D - AEC_1$, and that term grows as challengers improve (the saving available from adopting gets bigger). Waiting pays only while each generation's improvement exceeds i times the saving being forgone, here, only while a robot generation cuts more than about $0.10 \times 45{,}032 \approx 4{,}500$ SAR per year off the AEC. When generational gains shrink below that bar, the treadmill stops and the firm buys. "There will always be a better model next year" is an argument with an expiry date, and Equation (19.1) computes it.

Learning curve. The empirical regularity that the unit cost of a repeated task falls by a fixed fraction each time cumulative count doubles: $C_n = C_1\, n^{\log_2 \varphi}$, where φ is the learning rate.

Learning rate. The fraction φ to which unit cost falls at each doubling: $\varphi = 0.85$ means each doubling leaves 85% of the unit cost, a 15% saving per doubling.

In practice. Estimate φ from your own first deployments, not the vendor's brochure. The vendor's 80% curve was measured on their crews integrating their product into cooperative sites. Your rate carries your plant's permitting, your network, your union agreement. Two or three of your own data points on log-log paper beat any cited industry figure.

19.6 *The Learning Curve*

The first deployment is the most expensive one the firm will ever do. Commissioning the second robot goes faster: the survey template exists, the interlock design is reusable, the crew has done it before. This regularity, observed first in aircraft assembly in the 1930s and since in nearly every repeated industrial task, has a standard model. If the first unit costs C_1 and the *learning rate* is φ, the n-th unit costs

$$C_n = C_1\, n^b, \qquad b = \log_2 \varphi = \frac{\ln \varphi}{\ln 2}, \tag{19.2}$$

so that every *doubling* of cumulative count multiplies the unit cost by φ: $C_2 = \varphi\, C_1$, $C_4 = \varphi^2 C_1$, $C_8 = \varphi^3 C_1$. The exponent b is negative (for $\varphi = 0.85$, $b = \log_2 0.85 = -0.2345$), and the curve is a straight line on log-log paper, which is how φ is estimated from records.

Example 19.5— Budgeting eight deployments

Najd Copper plans to roll the inspection robot out across eight plant areas. The first integration cost 140,000 SAR (Example 19.1), and the engineering team estimates an 85% learning rate for the integration work (the hardware price is the same for every unit and is budgeted separately). (a) Tabulate the integration cost of each deployment and the cumulative total. (b) Compare the total against the two naive budgets: eight at first-unit cost, and eight at the mature eighth-unit cost. (c) From which deployment onward does an integration cost less than 100,000 SAR?

Solution.

Step (a): the unit-cost table.

With $b = \log_2 0.85 = -0.2345$, Equation (19.2) gives $C_n = 140{,}000\, n^{-0.2345}$:

deployment n	n^b	unit cost C_n (SAR)	cumulative (SAR)
1	1.0000	140,000	140,000
2	0.8500	119,000	259,000
3	0.7729	108,208	367,208
4	0.7225	101,150	468,358
5	0.6857	95,994	564,352
6	0.6570	91,977	656,329
7	0.6337	88,712	745,041
8	0.6141	85,977	831,018

Table 19.3. The learning-curve budget for eight robot integrations at $\varphi = 0.85$. Doublings, $n = 1 \to 2 \to 4 \to 8$, multiply the unit cost by 0.85 each time; the highlighted row is the first deployment under 100,000 SAR.

Check the doublings: $C_2/C_1 = 0.85$, $C_4/C_2 = 101{,}150/119{,}000 = 0.85$, $C_8/C_4 = 85{,}977/101{,}150 = 0.85$. The total is $\boxed{831{,}018 \text{ SAR}}$, an average of about 103,877 SAR per deployment.

Step (b): against the naive budgets.

Eight at first-unit cost: $8 \times 140{,}000 = 1{,}120{,}000$ SAR, an overstatement of about 289,000 SAR (35% too high), which pads the proposal and invites a budget cut later. Eight at eighth-unit cost: $8 \times 85{,}977 = 687{,}816$ SAR, an understatement of about 143,000 SAR (17% too low), which is how scale-ups run out of money in the middle. The learning curve is the honest line between the two.

Step (c): the threshold deployment.

Solve $C_n < 100{,}000$: from the table, $C_4 = 101{,}150 > 100{,}000 > C_5 = 95{,}994$, so the threshold is crossed at $\boxed{n = 5}$. (Analytically, $n = (100/140)^{1/b} = 4.2$, confirming the fifth deployment.)

Insight. The learning curve's clock is the doubling, not the year. Going from unit 1 to unit 2 saves 15%; going from unit 4 to unit 8, four deployments of work, saves the same 15%. Early experience is cheap to buy and pays the steepest part of the curve, which is one more economic argument for the pilot of Section 19.8: the pilot is not only information, it is position on the curve.

Figure 19.3 plots the eight costs of Table 19.3; compare the vertical drop at each doubling, and how much work each one costs to reach.

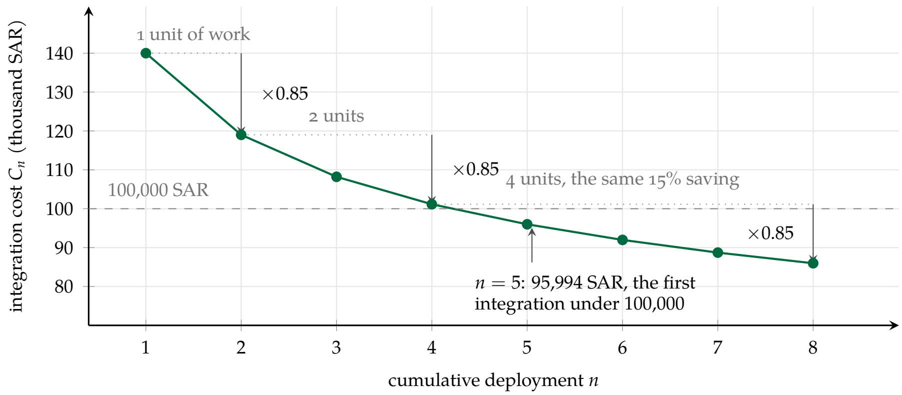

Figure 19.3. The learning curve for the eight robot integrations of Example 19.5 at an 85% learning rate. Each doubling of the cumulative count, 1 to 2 to 4 to 8, cuts the unit cost by the same 15%, so the fourth doubling buys as much saving as the first while taking four times the work.

19.7 *Labor: Reallocation and Displacement*

The most sensitive line in any automation case is the people. The economics here must be handled exactly, because both advocacy and opposition routinely misprice it. The rule is the same incremental principle as everywhere else: count the cash flows that actually change, and nothing else. Two honest cases cover practice.

- **Reallocation.** The displaced hours move to work the firm would otherwise *pay for*: a contracted maintenance backlog, purchased outside services, vacancies that would have been filled. Salaries continue unchanged, so the benefit is *not* the salary; it is the avoided external spending the redeployed hours now cover. If no such work exists, the reallocation benefit is zero, however busy everyone looks.

- **Displacement.** The positions are eliminated. The benefit is the salaries and burdens that actually stop, and the cost side gains a real cash flow: severance, notice periods, and any retraining owed. These are future consequences of the decision, so, unlike sunk costs, they belong in the analysis.

A single analysis may not claim both: hours cannot be simultaneously redeployed to the backlog and removed from the payroll. Claiming the salary *and* the covered contractor work is the double-count this chapter's closing audit problems are built to catch.

Reallocation vs. displacement. Reallocation redeploys freed hours to work the firm would otherwise purchase; its value is the avoided purchase. Displacement removes the positions; its value is the stopped salaries net of severance. An analysis may claim one or the other, never both.

Example 19.6 — Two honest labor cases

The inspection robot frees the two Najd Copper technicians ($2 \times 48{,}000 = 96{,}000$ SAR per year). The maintenance superintendent currently pays outside contractors 110,000 SAR per year for a backlog the two could cover. Alternatively, the positions could be eliminated with a one-time severance of 75,000 SAR. At $i = 10\%$ over the robot's 5-year life, value each case and recommend.

Solution.

Case R: reallocate.

Salaries continue in both futures, so they cancel; the change is the contractor line:

$$B_R = \boxed{110{,}000 \text{ SAR per year}}.$$

Case S: displace.

The salaries stop; the severance is a year-0 outflow annualized over the study period:

$$\begin{aligned} B_S &= 96{,}000 - 75{,}000\,(A/P,\,10\%,\,5) = 96{,}000 - 75{,}000\,(0.26380) \\ &= 96{,}000 - 19{,}785 = \boxed{76{,}215 \text{ SAR per year}}. \end{aligned}$$

Recommendation.

Reallocation is worth $110{,}000 - 76{,}215 = 33{,}785$ SAR per year more than displacement, *and* it keeps two trained technicians in a plant that struggles to hire them. We recommend reallocation, with one condition attached in writing: the contractor spending must actually stop. If the backlog were fiction, Case R would be worth zero and Case S would be the only real option, which is why the superintendent's contractor invoices, not the org chart, are the evidence an auditor asks for.

The analysis is deliberately neutral: this book prices alternatives, it does not choose sides in labor policy. What the method insists on is honesty in both directions: no salary counted while the person still draws it, and no severance hidden when it will in fact be paid.

19.8 Pilot, Then Scale

Nothing in the inspection-robot case is known with certainty: the duty ceiling, the detectable-failure fraction, the learning rate, the labor outcome. chapter 12 would attack this with sensitivity analysis, and the utilization table in the next section does. But automation offers something better than analysis: the decision can be *staged*. Buy one robot, run it for a year, measure the true utilization and the true integration cost, and let the eight-site rollout decision be made by the measurements. Staging converts a 400,000-SAR leap into a small experiment plus an informed choice, and the machinery for valuing that structure exists in this book: the decision tree of chapter 13 rolls back the pilot's outcomes, and chapter 15 prices the option to expand or abandon that the pilot creates. The learning curve adds its own argument: the pilot is deployment $n = 1$, the steepest lesson on the curve, purchased at the smallest scale.

19.9 The Inspection Robot, End to End

The chapter's tools now assemble into one recommendation. The pieces so far: installed basis 400,000 SAR (Example 19.1); honest benefit pool 174,000 SAR per year at full duty, built from the stoppage, labor, and safety lines (Example 19.3); reallocation chosen for the technicians (Example 19.6); one generation waited for and the improved model now installed at the original model's audited economics (Example 19.4). One estimate rules them all: *utilization*, the fraction of the planned inspection duty the robot actually achieves. Every benefit line scales with it, because uninspected galleries save no stoppages and un-relieved technicians cannot be reallocated.

Utilization. The achieved fraction of an automation asset's planned duty. Benefits scale with it; the vendor assumes 1.0; engineering ceilings (charging, maintenance windows, takt) usually cap it below 0.9.

Example 19.7 — The recommendation, with its sensitivity

Screen the inspection robot at $i = 10\%$: installed basis 400,000 SAR, salvage 40,000 SAR after 5 years, operating cost 30,000 SAR per year, benefit pool 174,000 SAR per year scaled by utilization u. The engineering ceiling is

$u = 0.85$ (charging and maintenance windows); the pilot measured $u = 0.78$ over its last two quarters. Find the break-even utilization and state the recommendation.

Solution.

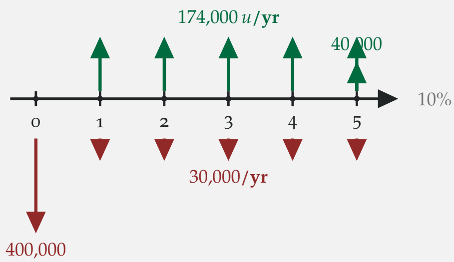

Figure 19.4. The inspection robot end to end: installed basis, five years of utilization-scaled benefits against operating cost, and salvage. The single symbol u carries the whole decision.

Step 1: the robot's annual equivalent cost.

From Example 19.3, $AEC = 128{,}968$ SAR per year.

Step 2: break-even utilization.

Set benefits equal to cost, the break-even move of Section 12.4:

$$u^* = \frac{128{,}968}{174{,}000} = \boxed{0.741},$$

so the robot must deliver at least 74.1% of its planned duty to pay its way.

Step 3: the sensitivity table.

Net annual advantage $= 174{,}000\,u - 128{,}968$:

utilization u	benefits (SAR/yr)	net (SAR/yr)
0.70	121,800	−7,168
0.741 (break-even)	128,968	0
0.78 (pilot)	135,720	+6,752
0.85 (ceiling)	147,900	+18,932
1.00 (vendor)	174,000	+45,032

Table 19.4. Utilization sensitivity for the inspection robot. The project lives in the narrow band between the pilot's measurement and the engineering ceiling; the vendor's row is not achievable.

Step 4: recommendation.

We recommend that Najd Copper adopt the robot, because the pilot's measured utilization of 0.78 clears the break-even of 0.741 and yields about

+6,752 SAR per year, with upside to +18,932 at the engineering ceiling, and with two conditions: the contractor backlog reallocation of Example 19.6 is confirmed in the maintenance budget, and utilization is reported quarterly, because a slip of eight points below the pilot's measurement erases the margin. The vendor's row in the table is recorded for reference and for chapter 20's files: it assumed a duty the machine cannot physically keep.

19.10 *Case Study from the Literature*

Published techno-economic studies of robots apply exactly this chapter's procedure: define the incumbent process as the baseline, price the robot scenario, convert the differences to cash flows, and screen with the standard indicators. The case below reads one such study, an open-access cost–benefit analysis of a facade-installation construction robot, with the chapter's eyes.

Reference. Hu, Iturralde, Linner, Zhao, Pan, Pracucci, and Bock (2021), *Buildings*, 11(1):8. Open access; full reference at the end of the chapter.

Example 19.8 — Case study: a facade-installation robot in the UK

Hu et al. (2021) evaluate the Hephaestus cable-driven robot, a system that installs curtain-wall modules on building facades, against the conventional manual installation method, from the viewpoint of a UK facade-installation company that would own the robot for its estimated five-year life. The study design fixes the workload so both scenarios deliver the same revenue; the benefits are therefore the avoided costs of the conventional method, and the costs are the money spent to run the robot, this chapter's service-project framing exactly. Key published figures: the robot scenario is 2.2 times as productive as the manual method; adopting it requires a one-month training program (12,000 EUR of fees, 3,000 per worker, plus 9,900 EUR of salaries paid during training), after which one worker becomes a specialist operator earning 30% more; the resulting benefit–cost ratio is 1.95, and the payback period is 21.21 months. A break-even analysis finds the robot worthwhile wherever construction labor costs exceed about 3.45 EUR per hour (2020 prices), which covers most G20 countries. The sensitivity analysis ranks the productivity coefficient first, then the labor hourly rate, then the robot system cost, then crane rental, then the annual wage increase. (a) Map the study onto this chapter's framework. (b) Interpret the payback and BCR against the five-year life. (c) Read the sensitivity ranking and the break-even wage with the chapter's tools.

Solution.

Step (a): the mapping.

The manual method is the defender with zero market value; the robot is the challenger. The equal-revenue design is the honest conversion of Section 19.3 enforced by construction: benefits are only the avoided conventional costs, so no unsold capacity can leak into the numerator. The training package, fees *and* the salaries paid while training, plus the operator's permanent 30% premium, is the retraining component of the installed basis and operating costs (Section 19.2): the study prices not just the course but the downtime and the lasting wage structure, which many vendor sheets omit. The five-year horizon is the challenger's assumed service life, chosen because the authors judge a robot with a payback beyond five years unattractive to an individual firm.

Step (b): the indicators.

A 21.21-month payback against a 60-month life means the robot spends roughly a third of its life repaying itself and the rest earning, consistent with a benefit–cost ratio of 1.95: discounted savings about double the discounted spending. Note the contrast with Example 19.3: this 21-month figure is what a defensible payback claim looks like, built on an equal-workload baseline with training and downtime priced, not a brochure's 5.5 months.

Step (c): the ranking and the break-even.

The sensitivity ranking puts the productivity coefficient (2.2) first, the study's version of this chapter's utilization: the single scaling factor on the benefit side dominates everything on the cost side, exactly as in Table 19.4. The break-even wage is the break-even move of Section 12.4 solved on a different variable: instead of asking what duty the robot must achieve, the authors ask what labor price the defender must charge for the challenger to win, about 3.45 EUR per hour, and then read the answer as a map of where in the world the robot competes. High-wage markets adopt first; as robot costs fall and wages rise, the frontier moves, which is Section 19.5's obsolescence current seen from the vendor's side.

Insight. The study's most transferable idea is the controlled comparison: fix the workload, let only the method vary, and benefits can only be avoided costs. Any automation case that cannot survive being restated this way, same output, what spending changes?, is counting something it should not.

19.11 *Chapter Summary*

- An automation or AI proposal is an ordinary **challenger**: compare annual equivalent costs against the incumbent process (the defender, usually with zero market value) under the rules of chapter 16. No new economics is created by new technology.
- The challenger's **installed cost basis** is quote plus **integration** plus **retraining**; on brownfield sites the integration line routinely rivals or exceeds the vendor's quotation. Sunk studies and pilots stay out.
- Benefits are converted **honestly**: cycle time is worth the spending it displaces, scrap at the measured delta, safety at the ledgered exposure costs, and 24/7 availability only against real demand. Write explicit zeros for benefits that fail the test.
- **Vendor ROI** fails at both ends of the fraction: bare-quote denominator, full-utilization brochure numerator. Audit by rebuilding the basis, re-converting the benefits, and restating at the MARR, then judge the project on the honest numbers.
- With **fast challenger obsolescence**, waiting one year pays exactly when the capitalized improvement beats the forgone saving: $(AEC_1 - AEC_2)/i > AEC_D - AEC_1$ (Equation (19.1)). The rule also ends the treadmill: buy when generational gains fall below i times the forgone saving.
- The **learning curve** $C_n = C_1\, n^{\log_2 \varphi}$ cuts unit cost by the factor φ at each *doubling* of count. Budget scale-ups on the curve, not at first-unit cost (too high) or mature cost (too low).
- **Labor** is priced neutrally: reallocation is worth the avoided external spending; displacement is worth the stopped salaries net of severance; never both at once.

- **Pilot-then-scale** stages the decision: the pilot buys the utilization measurement, the first position on the learning curve, and the option to expand or abandon (chapter 13, chapter 15). **Utilization** is usually the estimate the whole verdict turns on.

19.12 *Exercises*

The problems build in the chapter's order: first the installed cost basis and the honest conversion of claimed benefits, then quick checks on the learning curve, then longer problems combining break-even utilization, learning-curve budgets, and the adopt-now-or-wait rule. The two audits that close the set are a vendor sheet and an analyst's memo whose flaws are a bare-quote denominator, an unreachable utilization, and double-counted labor.

Warm-up and lecture practice

Exercise 19.1. Wadi Phosphate is quoted 150,000 SAR for a warehouse AGV (automated guided vehicle). The project file also shows: a site survey completed and paid last month, 12,000 SAR; floor marking and charging bays, 38,000 SAR; integration with the warehouse management system, 45,000 SAR; operator retraining, 17,000 SAR; an initial spare-parts kit purchased with the vehicle, 10,000 SAR. Determine the installed cost basis, stating what is excluded and why.

Exercise 19.2. Najd Copper considers a survey drone for its stockpiles: installed cost 95,000 SAR, salvage 15,000 SAR after 4 years, operating cost 9,000 SAR per year, $i = 10\%$. Claimed benefits: it replaces a manned stockpile survey contracted at 7,500 SAR per quarter; faster inventory counts are expected (from a measured trial) to cut concentrate write-offs by 8,000 SAR per year; and it ends technicians' walking on stockpiles, whose permit and safety overhead is 4,000 SAR per year. Convert the benefits, compute the AEC with $(A/P, 10\%, 4) = 0.31547$, and decide.

Quick checks (multiple choice)

Exercise 19.3. A task's first unit costs 20,000 SAR and the learning rate is 80%. The cost of the *eighth* unit is nearest to:

(A) $20{,}000 \times 0.80 = 16{,}000$ SAR
(B) $20{,}000 \times 0.80^3 = 10{,}240$ SAR
(C) $20{,}000 \times 0.80^8 = 3{,}355$ SAR
(D) $20{,}000/8 = 2{,}500$ SAR

Exercise 19.4. Which of the following belongs in an automation challenger's installed cost basis?

(A) The consultant study, completed last year, that recommended automating.
(B) The revenue of the product line being automated.
(C) The commissioning and safety-interlock work needed before the robot can run.
(D) The salaries of the operators during the years after adoption.

Exam-style problems

Exercise 19.5. Predictive maintenance on Wadi Lithium's brine pumps. A vendor offers an AI predictive-maintenance system: sensors and license quoted at 96,000 SAR. Installation and OT-network integration will cost 54,000 SAR, and configuring the models with staff training 20,000 SAR. A trial sensor kit bought last year for 15,000 SAR is now obsolete. System life 4 years, salvage 10,000 SAR, subscription and upkeep 18,000 SAR per year, $i = 10\%$, $(A/P, 10\%, 4) = 0.31547$. The pump fleet suffers 4 unplanned failures per year at 30,000 SAR each; the reliability engineer judges 60% of failure modes predictable from vibration and temperature. Maintenance overtime of 10,000 SAR per year and parts expediting of 8,000 SAR per year would also stop. The vendor's sheet additionally claims 40,000 SAR per year of "maintenance labor productivity," but the crews are fixed and fully occupied. (a) (*6 pts*) Determine the installed cost basis, stating any exclusions. (b) (*6 pts*) Compute the system's annual equivalent cost. (c) (*6 pts*) Convert the benefits honestly, with an explicit zero where deserved. (d) (*4 pts*) The sensors as quoted cover only a fraction u of the historical failure modes, and all benefits scale with u. Find the break-even coverage u^*. (e) (*3 pts*) The failure log shows the quoted sensor set covers $u = 0.70$. Decide, in one sentence.

Exercise 19.6. An AGV fleet and one year of patience at Najd Copper. The warehouse plans 4 AGVs. Commissioning the first costs 60,000 SAR and the learning rate is 90% ($b = \log_2 0.90 = -0.152$); hardware is 130,000 SAR per vehicle, unaffected by learning. (a) (*8 pts*) Tabulate the four commissioning costs and total the fleet budget (hardware plus commissioning). (b) (*4 pts*) By how much would budgeting all four commissionings at first-unit cost overstate? (c) (*8 pts*) Adopting the fleet now saves a net 52,000 SAR per year over the manual tugger operation. Next year's model would cut the fleet's annual equivalent cost by 4,000 SAR per year. Using Equation (19.1) at $i = 10\%$, should the firm adopt now or wait a year? (d) (*5 pts*) At what annual-AEC improvement would waiting start to pay?

Assessing outside recommendations

Automation proposals usually arrive as someone else's numbers, a vendor's sheet, a consultant's memo, an enthusiastic AI-generated analysis. The two problems below are audits: decide whether to accept the recommendation, name every flaw, and correct the numbers. The flawed work is quoted exactly as received.

Exercise 19.7. Wadi Phosphate's bagging line receives this vendor proposal for a palletizing robot, at $i = 10\%$:

> *"Robot price: 240,000 SAR. The robot replaces one loader per shift across your three shifts (3 × 45,000 = 135,000 SAR per year). Payback:* $240{,}000/135{,}000 = 21$ *months, a 56% annual return. Approval is recommended without delay."*

The plant's own file shows: conveyor interfacing, guarding, and PLC work 85,000 SAR; retraining 15,000 SAR; the line runs only two shifts for half the year (so 2.5 shift-loaders displaced on average, not 3); the night loader also performs QC sampling that will still require half a position; robot life 6 years, salvage 30,000 SAR, operating cost 16,000 SAR per year, $(A/P, 10\%, 6) = 0.22961$. Accept or reject the memo? Name each flaw and correct the numbers.

Exercise 19.8. Najd Copper's warehouse automation memo, prepared by a junior analyst, reads:

> *"AGV system, installed 300,000 SAR, life 5 years, salvage 25,000 SAR, operating cost 20,000 SAR per year, MARR 10%. Benefits: the two forklift drivers (84,000 SAR per year together) will be redeployed to the maintenance backlog, saving the 100,000 SAR per year we currently pay contractors for that work; total benefit* $84{,}000 + 100{,}000 = 184{,}000$ *SAR per year. Net advantage after capital charges: about* $+89{,}000$ *SAR per year. Strongly recommended."*

With $(A/P, 10\%, 5) = 0.26380$, audit the memo. If the drivers were instead separated with a one-time severance of 50,000 SAR, would that change the recommendation?

References

Hu, R., Iturralde, K., Linner, T., Zhao, C., Pan, W., Pracucci, A., & Bock, T. (2021). A simple framework for the cost–benefit analysis of single-task construction robots based on a case study of a cable-driven facade installation robot. *Buildings*, 11(1), 8. https://www.mdpi.com/2075-5309/11/1/8

Park, C. S. (2012). *Fundamentals of Engineering Economics* (3rd ed.). Pearson. [Cited throughout as "Park, 3rd ed."]