

The Robotics Economy

Physical AI, Machine Labor and the Economics of Automation

Central question. Can robotics reduce the marginal cost of physical labor in the same way artificial intelligence is reducing the marginal cost of certain forms of intelligence?

Finding: partially supported. Machine labor is already cheaper than human labor for a narrow, structured class of physical tasks, and the cost of that class is falling. It is not yet cheap for general physical work, and the binding constraint below roughly five dollars per human-equivalent hour stops being price and becomes capability, integration capacity and institutional acceptance.

Report scope Twenty analytical sections · five original economic models · twenty-five exhibits · robotics market outlook · investment opportunity map · assumption appendix · contradiction register · thesis scorecard

Research cut-off 29 August 2026. All web sources accessed 27–29 August 2026.

Evidence discipline Every material quantitative claim is classified as observed data, company-reported result, announced target, third-party estimate, research assumption or author calculation.

Not investment advice The investment section identifies capital-allocation themes and diligence gates. It contains no security recommendations, rankings, valuations, price targets or return expectations.

RESEARCH METHOD

How this report was built

This report treats robotics as an economic transformation of labor and capital rather than as a technology trend. The organising chain runs: AI capability → Physical AI → robot capability → robot production → hardware and deployment costs → productive utilization → effective robot labor cost → human-robot cost parity → automation adoption → productivity → labor markets → production costs → output and prices → macroeconomic consequences. Each link is tested separately, and several of them fail under current evidence.

Research cut-off

29 August 2026. Statistical sources are the most recent releases available on that date: IFR *World Robotics 2025* (reference year 2024, published 25 September and 7 October 2025), BLS Employer Costs for Employee Compensation for March 2026 (published 12 June 2026), BLS Occupational Employment and Wage Statistics for May 2025, A3 order data through the second quarter of 2026, BloombergNEF's December 2025 battery price survey and IEA *Electricity 2026*. Company evidence runs through August 2026.

Stages completed before writing

1. **Source discovery.** Primary statistical agencies and standard-setting bodies first (IFR, BLS, OECD, IMF, UN DESA, World Bank, IEA), then peer-reviewed economics, then company filings and primary company communications, then journalism only where no primary record exists.
2. **Evidence collection.** Every quantitative claim was taken from the issuing organisation's own release wherever one exists, with the reference year recorded separately from the publication year.
3. **Definition and unit validation.** Six distinctions are enforced throughout and are the main source of confusion in public robotics commentary: purchase price versus total cost of ownership; scheduled versus operating versus productive hours; productive robot hour versus human-equivalent productive hour; produced versus shipped versus deployed versus operating productively; prototype versus pilot versus commercial deployment; and task automation versus job automation.
4. **Contradiction checking.** Where sources disagree, both are shown with the reason for the disagreement. Eleven material contradictions are registered in Appendix D rather than resolved by preference.
5. **Economic modelling.** Five original models were specified and computed from explicit assumptions. All model code is deterministic and the assumption set is published in Appendix A.
6. **Sensitivity analysis.** Eleven input variables were swept independently and in combination to identify which ones actually move the answer.
7. **Thesis testing.** Eight sub-claims were scored against the evidence in Appendix E. The overall verdict is derived from those scores, not asserted in advance.
8. **Report writing.**
9. **Citation and numerical-claim audit.** Every numeric statement was re-checked against its source, every model output re-derived from the published assumptions, every exhibit re-checked for unit and label consistency, and every claimed date verified.

Evidence taxonomy

Important evidence carries one of six labels. Nothing in this report presents an announced production target as actual production.

Label	Meaning
OBSERVED DATA	Measured and published by a statistical agency, regulator or independent counting body.
COMPANY-REPORTED	Stated by the company about its own operations; not independently verified.
ANNOUNCED TARGET	A plan, forecast or capacity ambition. Never treated as output.
THIRD-PARTY ESTIMATE	Produced by a research house, bank or consultancy from a proprietary method.
RESEARCH ASSUMPTION	An input this report chose to expose sensitivity. Not an estimate of the world.
AUTHOR CALCULATION	Derived here from labelled inputs using a published formula.

What this report deliberately does not do

It does not aggregate industrial and service robot counts into a single "robot" total, because the IFR itself warns that the service-robot series is a producer sample whose composition changes each year.^[3] It does not apply the learning rate of batteries, solar modules or semiconductors to robots by analogy. It does not convert a bill-of-materials figure into a labor cost. It does not forecast; the scenarios in Sections 10 and 12 are scenarios, and are labelled as such wherever they appear.

CONTENTS

FRONT MATTER

Research method — how this report was built

1. Executive summary — twelve research conclusions

Key numbers

PART I — FROM INTELLIGENCE TO EMBODIMENT

2. From generative AI to Physical AI — does LLM-style scaling transfer to control?

PART II — THE ROBOTICS ECONOMY TODAY

3. The global robot economy — installations, stock, density, geography, segments

4. Where robots actually work today — task-level automation readiness

PART III — THE ECONOMICS OF A ROBOT

5. What does a robot actually cost? — price versus total cost of ownership

6. The robot cost curve — what falls, what does not, and why

7. Robot unit economics — Model 1: cost per productive hour

PART IV — ROBOTS VERSUS HUMAN LABOR

8. Robots versus human labor — the two cost stacks, normalised

9. Human-robot cost parity — Model 2: the break-even equation

10. The \$10-per-hour robot — Model 3: the automation frontier

PART V — PRODUCTIVITY ECONOMICS

11. Robotics and productivity — reconciling contradictory literatures

12. The 24/7 economy — Model 4: utilization economics

PART VI — DEMOGRAPHICS AND LABOR SCARCITY

13. Demographics and labor scarcity

14. Robots as labor-scarcity technology

PART VII — THE TWO SYSTEMS

15. China's robotics advantage

16. America's Physical AI advantage

PART VIII — THE FLYWHEEL

17. The robotics flywheel — a link-by-link failure audit

PART IX — MACROECONOMICS

18. The deflationary machine — cost decline is not price decline

19. Labor, wages and inequality

PART X — JUDGMENT

20. Conclusion — what becomes scarce?

Model 5: Robot cost sensitivity matrix

Robotics Market Outlook and Valuation — bottom-up sizing, 2026–2030, with 2035 and 2040 scenarios

Investment Opportunity Map

APPENDICES

A. Model definitions and assumption appendix

B. Automation-frontier task screen

- C. Human and robot cost normalisation tables
- D. Contradiction register
- E. Thesis scorecard
- F. Exhibit register
- Bibliography

EXECUTIVE SUMMARY

1. Twelve research conclusions

These are conclusions reached by the research, not summaries of the chapters that follow. Each is stated so that it could be falsified by evidence.

1 • Robot labor is already cheaper than human labor — for a narrow class of tasks, and the class is defined by structure, not by sophistication

A specialised industrial cell at a \$75,000 platform price, run on two shifts with a 0.92 uptime and 0.95 task-success factor, prices at **\$11.75 per productive robot hour** and **\$7.83 per human-equivalent hour** once its 1.5× throughput advantage is applied [AUTHOR CALCULATION](#). Fully loaded US warehouse labor is \$29.11 per hour [AUTHOR CALCULATION](#) from BLS wage and compensation data.^{[9][10]} Parity in that class was reached years ago, which is why 4.66 million industrial robots are already operating.^[1] The interesting question is not whether robots can be cheap; it is which tasks let them be.

2 • The general-purpose platform is not yet at parity, and the gap is larger than list prices suggest

The same \$75,000 price in a general-purpose configuration — lower integration, shorter life, higher maintenance, 0.85 uptime, 0.75 task success and 0.70 relative productivity — prices at **\$24.54 per human-equivalent hour** [AUTHOR CALCULATION](#). That is below the \$29.11 US benchmark in nominal terms but fails a normal 30 per cent switching margin. On the only multi-month operating record a manufacturer has published, the figure is **\$84.13**.

3 • Public evidence contains exactly one usable humanoid utilization observation, and it is 58 per cent

BMW reports that a Figure 02 unit at Plant Spartanburg logged **approximately 1,250 operating hours** across a ten-month deployment on ten-hour weekday shifts, moving more than 90,000 components in support of more than 30,000 X3 vehicles [OBSERVED DATA](#) [COMPANY-REPORTED](#).^{[25][26]} Against roughly 2,150 scheduled hours, that is an operating ratio near **0.58** [AUTHOR CALCULATION](#). Vendor and analyst cost models routinely assume 0.90 or higher. The difference between those two numbers is worth more than any plausible hardware price reduction.

4 • Utilization and relative productivity dominate hardware price in the cost equation

Across the ranges tested in Model 5, relative productivity moves human-equivalent cost by **\$30.67 per hour** and the scheduled operating window by **\$26.43**, against **\$26.27** for a purchase price swept from \$150,000 to \$20,000 and only **\$4.35** for financing rate [AUTHOR CALCULATION](#). A robot that costs half as much and works the same way saves less than a robot that costs the same and works twice as usefully.

5 • No single lever reaches \$10 per hour; four simultaneous ones do

From the \$24.54 base, cutting hardware to \$30,000 alone gives \$15.45; software improvements alone give \$15.28; a 22-hour window alone gives \$19.13; near-free supervision alone gives \$21.32. Combining hardware and software gives \$9.95; adding the longer window gives \$8.24; adding supervision gives \$5.74 [AUTHOR CALCULATION](#). The \$10 robot is a compound outcome, and each component of the compound has its own failure mode.

6 • Below about \$5 per hour the binding constraint stops being cost

In the 25-task frontier screen, 2 of 25 tasks clear a 30 per cent switching margin at a \$20 platform cost, 9 at \$10, 16 at \$5, and 18 at both \$2 and \$1 [AUTHOR CALCULATION](#). The count stops rising because five tasks that are economically viable at a dollar an hour remain blocked by dexterity, liability, regulation or trust. Cheap machine labor does not buy its way past a capability gate.

7 • Duty cycle, not wage level, explains why household and hospitality robotics lag

A task that can absorb work for only ten per cent of a robot's operating window multiplies its effective hourly cost tenfold. Household tidying needs a reference platform cost of roughly **\$0.66 per human-equivalent hour** to clear a switching margin against unpaid or low-paid domestic work; floor scrubbing in a large open building needs **\$8.00** AUTHOR CALCULATION. This is an arithmetic result, and it explains the observed order in which service robotics has commercialised better than any capability ranking.

8 • The productivity evidence is genuinely contradictory, and the contradiction is informative

Graetz and Michaels find robot densification added about **0.36 percentage points** to annual labor productivity growth across 17 countries from 1993 to 2007, roughly 15 per cent of aggregate productivity growth, with no significant effect on total hours.^[13] Acemoglu and Restrepo find that one more robot per thousand workers reduced the US employment-to-population ratio by **0.18-0.34 percentage points** and wages by **0.25-0.5 per cent** in exposed commuting zones.^[12] Aghion and co-authors find French firm-level employment *elasticities of +0.2 rising to +0.4* after automation.^[15] Dauth and co-authors find German manufacturing displacement fully offset by service employment.^[14] These are compatible: different units of observation, different exposure measures and different general-equilibrium channels. They are not compatible with any single confident forecast.

9 • Robot deployment and measured manufacturing productivity have moved in opposite directions in the United States

The world's operational robot stock nearly tripled between 2015 and 2024, from 1.63m to 4.66m, and the United States installed 34,000-40,000 industrial robots a year through the period. Over 2010-2022 US manufacturing labor productivity growth ran at **-0.5 per cent a year** against +3.4 per cent from 1987 to 2007, and manufacturing TFP growth fell from 1.4 per cent to 0.1 per cent OBSERVED DATA.^[11] The New York Fed calls this a puzzle and notes it is pervasive across industries and firm sizes. Any thesis that robots mechanically deliver measured productivity must explain this twelve-year counterexample.

10 • China's advantage is deployment and components; its density statistic was overstated and has been revised

China installed **295,000** industrial robots in 2024, 54 per cent of the world total, and holds just over **2 million** in operation; domestic suppliers took 57 per cent of Chinese installations, up from 47 per cent in 2023 OBSERVED DATA.^{[11][46]} Its robot density, however, was restated from 470 per 10,000 manufacturing employees for 2023 to **166** for 2024 after China's National Bureau of Statistics revised the manufacturing employment denominator.^{[21][4]} The robots did not disappear; the ratio was wrong. China also controls **91 per cent** of rare-earth refining and **94 per cent** of sintered permanent magnet production, which is the single most concentrated input in any robot's actuator.^[22]

11 • Humanoid shipments are real, small, and mostly not deployed

An independent research house counts roughly **19,100** humanoid units shipped globally in the first half of 2026, 97 per cent from Chinese manufacturers, with AgiBot at about 8,400 and Unitree at about 5,900 THIRD-PARTY ESTIMATE.^[33] The same source estimates roughly **250 Unitree units in active deployment** from more than 5,500 shipped in 2025. Shipped is not deployed; deployed is not operating productively. Tesla has publicly moved Optimus V3 to mid-2026 and declined to commit to a 2026 volume.^[35] Amazon withdrew its Blue Jay robotic workcell from operations in February 2026 after testing it in a single South Carolina facility.^[28]

12 • If machine labor scales, the binding constraints migrate to inputs that robotics does not control

Electricity is the clearest. The IEA expects global electricity demand to grow **3.6 per cent a year to 2030**, with data centres alone driving around half of US demand growth, against grid investment that must rise about 50 per cent from today's \$400 billion a year OBSERVED DATA.^[23] Rare-earth magnet capacity outside China meets under 20 per cent of projected 2035 demand.^[22] Integration and installation labor is scarce today and gets scarcer with volume. The constraint set moves from wage parity toward magnets, actuators, grid interconnection, installers and standards.

VERDICT ON THE CENTRAL QUESTION

Partially supported. Robotics is reducing the marginal cost of *selected* physical tasks, and Physical AI is widening that selection. But intelligence and physical labor scale differently in a way that is structural rather than temporary: a trained model can be copied at near-zero marginal cost, while every additional robot hour requires a machine that must be manufactured from constrained materials, powered, integrated into a specific workflow, maintained, insured and supervised. Robotics converts some physical labor from a scarce human input into a scalable capital service. It does not make physical labor free, and the evidence does not support the claim that it will.

Key numbers

Ten statistics that frame the report. Reference years are stated because several of these series were revised in the last twelve months.

542,076

Industrial robots installed worldwide in 2024 — the second-highest year on record, and the fourth consecutive year above 500,000

OBSERVED IFR World Robotics 2025^[1]

4.664m

Industrial robots in operation worldwide, end-2024, up 9% on 2023. The installed capital stock, not an annual flow

OBSERVED IFR^[1]

132

Global average robot density, robots per 10,000 manufacturing employees, 2024. Korea 1,220; Germany 449; Japan 446; US 307; China 166

OBSERVED IFR^[2]

54%

China's share of global industrial robot installations in 2024 (295,000 units); operational stock just above 2 million; domestic suppliers took 57% of the Chinese market

OBSERVED IFR^{[1][46]}

~200k

Professional service robots sold in 2024 from a 294-supplier sample: 102,900 transport & logistics, 42,000+ hospitality, 25,000+ cleaning, 19,500 agriculture, 16,700 medical

OBSERVED SAMPLE IFR^[3]

\$46.60

US private-industry employer cost per hour worked, March 2026: \$32.60 wages (69.9%) plus \$14.01 benefits (30.1%)

OBSERVED BLS ECEC^[9]

\$61,197

Average value per industrial robot ordered in North America in 2025 (36,766 units, \$2.25bn). H1 2026: \$64,796 — the unit value rose, it did not fall

AUTHOR CALCULATION from A3^{[42][43]}

1,250 h

Operating hours logged by a Figure 02 humanoid over a ten-month BMW Spartanburg pilot — about 58% of scheduled hours. 90,000+ components moved; 1.2m steps

COMPANY-REPORTED BMW^[25]

–16.6%

Fall in Japan's working-age population (15–64) from its 1995 peak of 87.5m to 72.9m in 2024. China's peaked in 2015, Germany's in 2016, Korea's in 2019

OBSERVED World Bank WDI^[21]

\$24.54

Modelled cost per human-equivalent hour for a general-purpose platform on base assumptions — versus \$7.83 for a specialised industrial cell and \$84.13 on pilot-observed utilization

AUTHOR CALCULATION Model 1

\$108

Global average lithium-ion battery pack price per kWh in 2025, down 8% year on year and 91% from 2010. China \$84; North America \$121

OBSERVED BloombergNEF^[24]

94%

China's share of global sintered permanent magnet production, 2024, up from about 50% in 2005; 91% of rare-earth refining; 60% of magnet-rare-earth mining

OBSERVED IEA^[22]

PART I — FROM INTELLIGENCE TO EMBODIMENT

2. From generative AI to Physical AI

Generative AI reduced the marginal cost of producing text, code and images. The economic question for robotics is whether the same mechanism — large models trained on large corpora, improving predictably with scale — can reduce the marginal cost of producing physical work. The mechanism partly transfers. The data economics do not.

2.1 What "Physical AI" actually names

Physical AI describes a closed loop that generative AI does not require: observe, decide, act, measure the consequence in the world, and recover when the consequence is wrong. A language model's error costs a token. A manipulator's error costs a part, a cycle, a machine, or a person. Every element of the technical stack below exists to manage that asymmetry.

Component	What it does	Economic significance, and what limits it
Embodied AI	Joins perception to a policy that emits motor commands, closing the sense-plan-act loop on real hardware.	Converts capability into throughput only if the loop closes fast enough and fails safely. Latency and safety are cost items, not features.
Vision-language-action models	Tokenise actions the way a language model tokenises words, so semantic knowledge from web-scale pre-training reaches control.	The main source of generalisation to unseen objects and instructions. Constrained by geometry, contact force and the fact that action tokens have physical units.
Robot foundation models	One model serving many tasks and, increasingly, many bodies.	Amortises training cost across a fleet. Value depends entirely on whether an update generalises — and on whether it can be validated safely before it ships.
Imitation learning	Learns from human demonstrations, giving a safe behavioural prior cheaply.	Demonstration collection is a linear human cost. It is the single largest hidden operating expense in current robot-learning programmes.
Reinforcement learning	Optimises a policy against a reward, including behaviours no human demonstrated.	Sample-expensive on real hardware; wear, breakage and safety bound how much real-world exploration is affordable.
Teleoperation	Humans drive the robot, supplying labels, rare recovery examples and a commercial fallback.	Makes a deployment work before autonomy does — but converts a capital cost into a variable labor cost, which is exactly what automation is supposed to remove.
Simulation & synthetic data	Generates variation, lighting, geometry and failure cases at near-zero marginal cost.	The cheapest scaling lever available. Bounded by the reality gap in contact dynamics, deformables and sensor noise.
Real-world robot data	The only data that contains true contact physics and true failure distributions.	Expensive, embodiment-specific and commercially guarded. This is where the analogy with language models breaks.
Edge inference	Runs the policy on the machine rather than in a datacentre.	Removes network latency and outage as operational risks; adds power, thermal and unit-cost burden to the robot itself.

2.2 Does LLM-style scaling transfer?

The evidence says: partly, and unevenly. Three published results define the current frontier.

Cross-embodiment transfer is real but modest. The Open X-Embodiment collaboration pooled **over one million robot trajectories** across **22 embodiments**, assembled from 60 datasets contributed by 34 labs, covering 527 skills [OBSERVED DATA](#).^[38] The resulting RT-1-X model outperformed the original task-specific methods by about **50 per cent** on small-scale domains, and RT-2-X showed roughly a **threefold** improvement on emergent skills transferred from another robot's data, across 3,600 evaluation trials on six robots.^[39] That is genuine positive transfer. It is also two orders of magnitude smaller, in raw scale, than any frontier language corpus, and the units are not fungible: a trajectory recorded on a bimanual static arm does not directly teach a legged mobile manipulator.

Generalisation to unseen environments now works, at a specific and instructive data scale. Physical Intelligence's $\pi_{0.5}$ was trained on roughly **400 hours** of mobile-manipulator data collected in about **100 home environments**, supplemented so heavily with other sources that **97.6 per cent** of first-phase training examples came from something other than the target robot [OBSERVED DATA](#).^[40] It executed 10–15 minute multi-stage cleaning tasks in three previously unseen homes, and performance improved monotonically as the number of distinct training locations rose through 3, 12, 22, 53, 82 and 104 — with the 104-location model roughly matching a model trained directly on the test home. That is a scaling curve in the number of *environments*, not in the number of tokens. It implies the binding input is environmental diversity, which is collected by people, in buildings, in real time.

Motion transfer across bodies is emerging. Google DeepMind reports that Gemini Robotics 1.5 transfers behaviours learned only on the ALOHA 2 platform to an Apptроник Apollo humanoid and to bimanual Franka arms without embodiment-specific specialisation, and that its embodied-reasoning sibling leads 15 academic spatial-understanding benchmarks [COMPANY-REPORTED](#).^[36] Its predecessor RT-2 raised performance on unseen scenarios from 32 to 62 per cent relative to RT-1 [COMPANY-REPORTED](#).^[37] These are research benchmarks. None of them is a factory service level, and none has been published as a paid-production intervention rate.

WHY LONG HORIZONS PUNISH SMALL ERROR RATES

A policy with 95 per cent per-step success completes a ten-step task without human intervention $0.95^{10} = 59.9$ per cent of the time. At 99 per cent per step it completes **90.4 per cent**. At 99.9 per cent, **99.0 per cent** AUTHOR CALCULATION. In the cost model of Section 7 this is the term Q , and the table below shows why a capability improvement that looks incremental in a benchmark is not incremental in a P&L. Moving per-step reliability from 95 to 99 per cent cuts the human intervention rate on a ten-step task by roughly three-quarters.

Per-step success	5-step task	10-step task	20-step task	50-step task
90.0%	59.0%	34.9%	12.2%	0.5%
95.0%	77.4%	59.9%	35.8%	7.7%
99.0%	95.1%	90.4%	81.8%	60.5%
99.9%	99.5%	99.0%	98.0%	95.1%

2.3 The six frictions that do not appear in language-model scaling

Friction	Why it is structural rather than temporary
Data scarcity	Robot episodes are produced in real time by real machines in real buildings. There is no equivalent of scraping the web. $\pi_{0.5}$'s 100 homes and 400 hours are close to the current state of the art for a single programme; a frontier text corpus is measured in trillions of tokens. ^[40]
Dexterity	Contact-rich manipulation of deformable, occluded or variable objects remains the hardest published problem. Rigid pick-and-place generalises; laundry, cable routing, food handling and soft-tissue contact do not, yet.
Latency	Control loops have hard deadlines. Cloud inference introduces network dependence into a safety loop, which pushes compute onto the machine — NVIDIA's Jetson AGX Thor delivers up to 2,070 FP4 TFLOPS with 128 GB in a 130 W envelope at \$3,499 for the developer kit, which is a real unit-cost and thermal burden. ^[41]
Safety	Every model update is a safety-relevant change to a machine that shares space with people. Validation is a fixed cost per release, and it does not fall with fleet size the way inference cost does.
Reliability	Failure consumes time, parts and equipment. The economically relevant metric is not benchmark success but mean time between human interventions in paid production — a number essentially no vendor publishes.
Out-of-distribution physics	The physical world's tail is not sampled by any corpus: a spill, a dropped pallet strap, a shifted rack, an unusual lighting condition, a part that arrives deformed. Simulation widens coverage but cannot enumerate it.

2.4 The economic conclusion of Part I

Software improvement raises the value of hardware that already exists. In the cost equation of Section 7 it operates on Q (task success), on φ (relative productivity) and on V (supervision), and Model 5 shows those three terms together dominate purchase price. That gives Physical AI a real option value: a fleet bought at today's prices becomes cheaper per unit of work without any change to its bill of materials. But the same equation shows why the transfer from generative AI is incomplete. Copying a model costs nothing. Copying a robot costs a robot.


PART II — THE ROBOTICS ECONOMY TODAY

3. The global robot economy

Industrial and service robot statistics are produced by different methods and must not be combined. Industrial robot counts use a stable definition and a maintained installed-stock model. Service robot counts are a survey of 294 suppliers that the IFR explicitly says is *not projected to the whole industry*, whose sample composition varies each year, and which it discourages comparing across report editions. ^[3] Adding the two produces a number with no meaning.

Exhibit 1 • Global industrial robot installations, 2014–2024, with the IFR outlook

Research question: is the annual flow of new industrial robots still growing? • **Units:** robot units installed per calendar year • **Period:** 2014–2024 observed; 2025 and 2028 outlook




Source: IFR, World Robotics 2022, 2023, 2024 and 2025 press releases. ^{[1][5][6][7]} **Evidence status:** **OBSERVED DATA** for 2014–2024; **ANNOUNCED TARGET** for the 2025 and 2028 outlook bars, which are IFR projections and are drawn unfilled for that reason.

Takeaway: the flow has plateaued, not accelerated. Installations more than doubled between 2014 and 2021 and have since moved sideways in a narrow band: 552,946 (2022), 541,302 (2023), 542,076 (2024). A thesis of exponential robot adoption is not supported by the installation series; what is growing is the stock, because installations exceed retirements.

Exhibit 2 • Global operational stock of industrial robots, 2015–2024

Research question: how large is the installed capital stock of machine labor, and how concentrated? • **Units:** robot units in operation at year end • **Period:** 2015–2024



Source: IFR World Robotics editions 2016–2025. ^{[1][5][6][7][46]} **Evidence status:** **OBSERVED DATA**. Values for 2016–2020 are the IFR published stock series; the China overlay is available only from 2021.

Takeaway: the stock nearly tripled in nine years, from 1.63m to 4.66m, and 43 per cent of it now sits in one country. Machine labor is already a large capital stock — the marginal question is not whether robots work but where the next unit earns its cost of capital.


3.1 Geography: one country, one segment, one decade

Market	2021	2022	2023	2024	2024 stock	Comment
World	517,385	552,946	541,302	542,076	4,664,000	Asia 74%, Europe 16%, Americas 9% of 2024 installs
China	268,195	290,144	276,288	295,000	>2,000,000	54% of world; domestic suppliers 57% of local market
Japan	47,182	50,435	46,106	44,500	450,500	Fourth consecutive annual decline
United States	34,987	—	37,587	34,200	—	68% of Americas installations; -9% in 2024
Rep. of Korea	31,083	—	31,444	30,600	—	Highest density in the world
Germany	23,777	—	28,355	26,982	—	32% of European installations; -5% in 2024
India	—	—	8,510	9,100	—	+7%; automotive 45% of installations
Europe total	—	—	92,393	85,000	—	EU 67,800, 80% of the European total

OBSERVED DATA. Source: IFR World Robotics 2022, 2024 and 2025. ^{[1][5][7]} Dashes indicate the figure was not published in the releases consulted; they are not zeros. Country installation counts are on the IFR's own definition of an industrial robot (ISO 8373) and are not comparable with national production statistics — see Section 15.

Exhibit 3 • Robot density and the 2024 methodology revision

Research question: how intensively is machine labor used relative to human manufacturing labor, and how reliable is the comparison? • **Units:** operational robots per 10,000 manufacturing employees • **Period:** reference years 2023 and 2024



Source: IFR, *Global robot density in factories doubled in seven years* (Nov 2024, reference year 2023) and *Robot density surges in Europe, Asia and the Americas* (2025, reference year 2024). ^{[2][4]} **Evidence status:** **OBSERVED DATA** on both vintages.

Takeaway: density is a ratio, and in 2025 the IFR restated the denominator. China moves from 470 to 166 not because robots were removed but because China's National Bureau of Statistics revised manufacturing employment upward. The global average falls from 162 to 132 for the same reason. Any analysis that treats robot density as a stable national "automation score" is comparing revised ratios across incompatible vintages. Registered as Contradiction C1.

3.2 The segments, kept separate

The table below is the report's segment map. Each row uses only figures published on that segment's own definition. The last column states, in one line, what the economics of that segment actually rest on — because the segments do not share a business model.

Segment	2024 volume	Where value comes from	Economic character
Industrial robots	542,076 installed 4.66m in operation	Repeatability, cycle time, uptime in an engineered cell	Mature capital goods. Cost parity settled; competition is on integration cost, service and financing.
Collaborative robots	~10.5% of industrial installs (2023); 19.6% of N. American units ordered (2025)	Low-volume, high-mix work; deployment without guarding	Lower price per unit (\$33,417 average North American order value in 2025) but higher share of value in application engineering.
Autonomous mobile robots & logistics	102,900 units (transport & logistics service robots, +14%)	Transport within a mapped, rule-governed building	The clearest RaaS market: IFR reports service-robot RaaS growing 31% overall and 42% in this segment.
Medical robots	~16,700 units (+91%)	Precision, ergonomics, asset utilization per system	Not labor substitution. Intuitive's da Vinci base reached 11,106 systems at end-2025 (+12%) with Q4 procedures +17%: the machine augments a clinician whose wage it does not replace.
Agricultural robots	~19,500 units (~6%)	Removing seasonal labor dependence; input savings	Duty cycle is the constraint, not capability. Seasonality is the reason — see Section 10.
Professional cleaning	25,000+ units (+34%)	Large, open, predictable floor area	Fastest-growing service segment. Structure of the environment, not sophistication of the robot.
Hospitality	42,000+ units (~11%)	Food running, delivery in constrained venues	Only service segment that shrank. Peak-hour demand means low duty cycle and poor unit economics.
Inspection, security, search & rescue	3,100 units (+19%)	Removing humans from hazardous or inaccessible zones	Value is risk avoidance and access, not wage substitution. Small volume, high willingness to pay.
Humanoid robots	~19,100 shipped in H1 2026 3RD-PARTY EST.	General-purpose form factor; brown-field deployment without re-engineering	Pre-commercial as labor. 97% of shipments Chinese; one third-party estimate puts Unitree's <i>active deployments</i> near 250 from 5,500+ units shipped in 2025.

Sources: IFR World Robotics 2025 industrial and service editions^{[1][3]}; IFR collaborative robots position paper^[8]; A3 North American order statistics^{[42][43]}; Intuitive Surgical Q4 2025 results^[47]; humanoid shipment estimates from Smart Analytics Global as reported August 2026^[33]. Volumes are on incompatible definitions by design and must not be summed.

CONSUMER SERVICE ROBOTS: THE LARGEST ROBOT POPULATION, AND THE ONE NOBODY COUNTS AS LABOR

The IFR reports close to **20 million** consumer service robots sold in 2024, up 11 per cent, overwhelmingly floor cleaners and lawn mowers **OBSERVED SAMPLE**.^[3] That is roughly forty times the industrial installation count. It is excluded from every labor-economics treatment of robotics, including this one, because these machines mostly substitute for unpaid household work rather than for waged labor. The exclusion is defensible for measuring wages and employment. It is a real omission for measuring welfare, and it is the single largest existing proof that a robot reaches mass volume when a task is structured, repetitive, low-consequence and tolerant of failure.

4. Where robots actually work today

The gap between demonstration and commercial deployment is where most robotics analysis goes wrong. This section separates the two using an explicit evidence ladder, then applies a task-level framework to nine sectors. The unit of automation is the task, not the job, and the determinant of value is how much of the environment can be engineered around the machine.

4.1 The commercial deployment evidence ladder

Six rungs, from a video to an audited operating result. Public claims about humanoid and general-purpose robotics cluster at rungs 1–3; the industrial robot industry lives at rungs 5–6. Nothing in this report is treated as evidence of labor substitution below rung 4.

Rung	Standard of evidence	Best current public example	What it does and does not prove
6	Audited multi-year operating economics	Industrial robot cells in automotive body shops; da Vinci systems (11,106 installed base, +12% in 2025, Q4 procedures +17%) ^[47] OBSERVED	Proves durable unit economics. Available only for mature categories, and even here procedure and cycle counts are reported, not cost per productive hour.
5	Fleet at scale in own operations	Amazon: one millionth robot deployed July 2025 across 300+ facilities; DeepFleet foundation model claimed to improve fleet travel time by 10% ^[27] COMPANY-REPORTED	Proves that co-designing building, inventory, software and machine works at enormous scale. Does not prove a robot can be dropped into someone else's building.
4	Paid multi-customer commercial deployment with volume metrics	Agility Robotics: 100,000+ totes moved for GXO; 65,000+ cumulative operating hours; 9 customer sites; named customers include GXO, Schaeffler, Toyota Motor Manufacturing Canada ^{[29][30][31]} COMPANY-REPORTED	Proves a bounded workflow is commercially viable. Tote handling is one loop, not an occupation. No intervention rate or cost per hour has been published.
3	Multi-month operating record inside a real production line	Figure 02 at BMW Spartanburg: ~1,250 operating hours over ten months on 10-hour weekday shifts; 90,000+ components; 1.2m steps; supported 30,000+ X3 vehicles ^[25] COMPANY-REPORTED OBSERVED BY THE CUSTOMER	The single most useful public data point in humanoid robotics, because BMW published hours as well as outputs. BMW calls it a pilot project. Implied operating ratio ~0.58.
2	Pilot in a single facility, outcome disclosed	Amazon Blue Jay: tested in one South Carolina facility, claimed to handle ~75% of item types and to consolidate three robotic stations into one — withdrawn from operations, disclosed February 2026 ^[28] COMPANY-REPORTED	The most instructive rung, because failures are usually not disclosed. A capable pilot from the most sophisticated operator in the industry did not survive contact with operations.
1	Demonstration, benchmark or announced target	Tesla Optimus: V3 reveal moved to mid-2026, Fremont line conversion, no 2026 volume commitment, ~10,000 unique parts per unit ^{[34][35]} ANNOUNCED TARGET	Proves engineering intent. Announced capacity is not output; a reveal is not a deployment; a benchmark is not a service level.

Exhibit 5. Research question: what standard of evidence supports each public claim of robot labor substitution? Units: qualitative ladder. Period: evidence published to 29 August 2026. Sources as cited per row. **Takeaway:** the humanoid category's best public evidence sits at rungs 3–4 and consists of hours and volumes, never cost per productive hour or intervention rate. The most economically decisive numbers are precisely the ones not disclosed.

4.2 Sector by sector: what the environment allows

Automotive. The founding case and still the densest. Welding, painting and heavy handling combine high repeatability, high throughput, expensive downtime and a plant that is rebuilt for each model cycle anyway — so re-engineering around the machine is not an incremental cost. Yet automotive is now a source of weakness rather than growth: A3 reports North American automotive OEM orders down **25 per cent** in the first half of 2026 while non-automotive customers took **56 per cent** of units.^[42] Saturation in the anchor sector is why installation totals have plateaued.

Electronics and semiconductors. Clean, controlled, precision-dominated and geographically concentrated: China hosted 64 per cent of global electronics-sector robot installations in 2024.^[46] In North America, semiconductors, electronics and photonics were the fastest-growing order category in H1 2026 at **+35 per cent** in units.^[42] Structure is engineered by necessity here, which is why robots work.

Warehouses and logistics. Floors, racks, barcodes, totes and routing rules can all be standardised, and the operator often owns the building. This is the only sector where a single firm's fleet has passed a million machines.^[27] It is also the sector where the distinction between transport and manipulation is sharpest: moving a tote is solved; picking an arbitrary item out of a mixed bin is not, and unloading a floor-loaded trailer is both unsolved and low duty cycle.

Agriculture. Viable niches exist in milking, guided vehicles, spraying and mechanical weeding. Delicate harvesting remains crop-specific and pre-commercial: the most advanced public strawberry system reports **~30 kg/hour** in field trials against a claimed 50 kg/hour potential, one operator supervising up to eight machines, and a ~4 per cent bruising rate, with commercial launch stated for April 2026 **COMPANY-REPORTED**, none of it independently verified.^[50] Note that the labor being replaced is cited by the same company at **\$43,000 per acre per year** for picking alone — a high-value target, which is precisely why the segment attracts capital despite a duty cycle of a few weeks a year.

Mining and inspection. The economic argument is not wage substitution but removing people from hazardous zones and reaching places people cannot. Willingness to pay is set by risk and access, not by the local wage, which is why these robots survive at low volumes and high prices.

Healthcare. Two distinct markets. Surgical robots augment a clinician and monetise precision and asset utilization — Intuitive's installed base grew 12 per cent to 11,106 systems in 2025 with fourth-quarter procedure volume up 17 per cent, and total revenue of \$10.06bn **COMPANY-REPORTED**.^[47] Logistics robots inside hospitals move supplies and linen and are ordinary AMRs. Patient handling — the task with the largest injury cost — is blocked by liability and trust, not by capability.

Hospitality, retail and cleaning. Cleaning grew 34 per cent in 2024; hospitality *fell* 11 per cent.^[31] The difference is duty cycle. A scrubber works a whole night shift on a predictable floor plan; a food-running robot works two peaks a day in a crowded room. Section 10 quantifies exactly this.

Construction. The geometry changes daily, multiple trades share the space, the site is weather-exposed and the work is sequenced by contracts rather than by a controller. Material transport and layout marking are tractable; finish trades in occupied space are not, and the obstacle is liability and coordination as much as manipulation.

4.3 Task-level automation readiness framework

The framework below scores twelve representative tasks on the seven dimensions that actually determine whether a robot earns its cost, plus a summed readiness verdict. Scores are 1 (favourable) to 5 (unfavourable) except where noted, and are author judgments calibrated to the deployment evidence in Section 4.1 **AUTHOR FRAMEWORK**.

Task	Fully loaded human cost/h	Robot capability	Deployment difficulty	Safety requirement	Utilization potential	Exception burden	Automation readiness
Spot welding, body-in-white	\$38.69	1	2	2	1	1	Commercial, mature. Fully automated for three decades.
Machine tending, CNC	\$30.78	1	2	2	1	2	Commercial. The core cobot application; constraint is integration cost, not capability.
Palletising / depalletising	\$29.11	1	2	2	2	2	Commercial. Mixed-SKU layer picking still needs vision tuning per site.
Tote and container movement	\$29.11	1	2	2	2	2	Commercial at fleet scale. Amazon and AMR vendors; the reference case.
Structured visual inspection	\$30.78	1	2	1	1	1	Commercial. Highest utilization of any task screened.
Floor scrubbing, open areas	\$27.46	1	2	2	3	2	Commercial. Fastest-growing service segment; duty cycle limits payback.
Each-picking, mixed bins	\$27.24	3	3	2	2	4	Near-commercial. Grasp success on the long tail of SKUs is the open problem.
Floor-loaded trailer unloading	\$29.11	3	3	3	4	4	Near-commercial. Capability is improving; arrival-pattern duty cycle is the harder problem.
On-site material transport	\$27.96	3	4	4	4	4	Pilot. Site access and changing geometry dominate.
Delicate fruit harvesting	\$20.87	4	4	2	5	5	Pre-commercial. Capability plus a duty cycle of weeks per year.
Patient lifting and transfer	\$28.85	4	4	5	5	4	Blocked. Liability, consent and trust bind before cost does.
General household tidying	\$17.21	5	5	4	5	5	Blocked. Worst score on every dimension; see Section 10.

Exhibit 4. Research question: which task attributes, not which sectors, determine automation readiness? Units: fully loaded employer cost in US dollars per hour; other columns 1–5 ordinal. Period: wages May 2025; readiness assessed August 2026. Sources: fully loaded cost is an **AUTHOR CALCULATION** from BLS OEWS May 2025 medians^[10] and the BLS ECEC compensation split^[9], per Appendix C; ordinal scores are **AUTHOR FRAMEWORK**. **Takeaway:** the commercial tasks are not the ones with the highest wages. They are the ones scoring 1–2 on utilization and exception burden. Willingness to pay follows structure, not wage level — which is the single most important pattern in the deployment record.

PART III — THE ECONOMICS OF A ROBOT

5. What does a robot actually cost?

The purchase price of a robot is roughly a quarter of what it costs to own and run one. Almost every public comparison of robot cost to human wages uses the wrong quarter. This section separates the two, and states plainly where public data is unreliable.

5.1 What the purchase price buys

Public humanoid bill-of-materials estimates are inconsistent, often promotional, and frequently derived from teardown speculation rather than procurement. This report therefore does not publish an "industry average" BOM. It publishes one company-reported anchor and an illustrative structure.

THE ONE DISCLOSED BILL OF MATERIALS

Agility Robotics disclosed a **\$125,000 bill of materials for Digit v4** in connection with its SPAC transaction, alongside a 35 lb payload, 5.5 ft reach and 16 hours of operation per charge **COMPANY-REPORTED**.^[29] Unitree's G1 lists from **\$13,500** **COMPANY-REPORTED**.^[32] These two numbers are nine times apart and describe different machines for different purposes. Neither is evidence about the cost of an hour of work, and a list price is not a bill of materials.

Subsystem	Illustrative share of BOM	What drives its cost, and where it is going
Actuators, motors, reducers, bearings	30–40%	The largest single block in any articulated machine. Harmonic and cycloidal precision reducers remain a concentrated, largely Japanese supply base; torque density depends on rare-earth permanent magnets. This is the subsystem least exposed to consumer-electronics deflation and most exposed to export controls.
Hands, grippers, end effectors	8–18%	Rises sharply with dexterity. Tactile sensing and multi-finger actuation add both cost and failure modes. Figure 03 adds tactile sensors and palm cameras to its hands. ^[26]
Sensors: cameras, depth, LiDAR, force/torque	8–15%	The one block reliably deflating, on consumer-electronics and automotive-ADAS volumes.
Compute and inference	6–12%	Jetson AGX Thor: 2,070 FP4 TFLOPS, 128 GB, 130 W, developer kit from \$3,499; 7.5× the compute of Orin at 3.5× the efficiency. ^[41] Performance per dollar improves fast; absolute cost and thermal load do not disappear.
Battery and power electronics	5–12%	Directly exposed to EV-scale deflation: global pack prices fell to \$108/kWh in 2025, –8% year on year, –13% in China. ^[24] But a humanoid carries 1–3 kWh, so a 10% pack price fall is a rounding error on a \$75,000 machine.
Structure, housings, wiring, thermal	10–18%	Falls with manufacturing scale, design-for-assembly and tooling amortisation. This is where volume genuinely helps.

RESEARCH ASSUMPTION for the share column: an illustrative structure used to show where cost sits, calibrated to the qualitative composition of articulated robots. It is not an estimate of any specific product's BOM and should not be quoted as one. Subsystem cost drivers are sourced as cited.

5.2 What the purchase price does not buy

Twelve cost categories sit outside the invoice. In the model of Section 7 they are the terms *I*, *M*, *S*, *O*, *E* and *V*, and together they are about three times the purchase price over a five-year life.

Category	Model term	Note
Software licences and inference	<i>S</i>	Increasingly a subscription; the vendor's recurring revenue is the buyer's recurring cost.
Fleet management and orchestration	<i>S</i>	Amazon's DeepFleet is a fleet-level model claimed to cut fleet travel time 10%. ^[27]
Integration and installation	<i>I</i>	Historically 1–2× robot hardware cost for an industrial cell; the general-purpose pitch is that this falls to 0.2–0.4×.
Factory or facility modification	<i>I</i>	Guarding, floor marking, charging infrastructure, network, lighting, fixture redesign.
Training, validation, safety engineering	<i>I</i> , <i>O</i>	A fixed cost per deployment and per material model update.
Maintenance, spares, repair logistics	<i>M</i>	Scales with duty cycle and with the number of actuated joints.
Downtime	<i>U</i> (denominator)	Does not appear as a cost line at all — it appears as fewer productive hours, which is worse.
Supervision and human intervention	<i>V</i>	The category most often set to zero in vendor models and most often non-zero in practice.
Energy	<i>E</i>	Small per hour, but a fleet-level and grid-level constraint at scale (Section 20).
Insurance and liability	<i>O</i>	Rises with human proximity; the reason patient handling is blocked.
Financing	<i>r</i>	Matters less than most buyers assume — Model 5 ranks it last of eleven variables.
Obsolescence	<i>n</i>	The reason an eight-year assumed life on a fast-moving platform is optimistic.


FIVE-YEAR TOTAL COST OF OWNERSHIP, GENERAL-PURPOSE BASE CASE AUTHOR CALCULATION

Line	Per year	Five-year total	Share of TCO
Purchase price, annualised at 10% over 5 years	\$19,785	\$98,924	34.2%
Integration investment (35% of price), annualised	\$6,925	\$34,624	12.0%
Maintenance and repair (12% of price p.a.)	\$9,000	\$45,000	15.6%
Software, fleet management, inference	\$9,000	\$45,000	15.6%
Insurance and other fixed	\$2,000	\$10,000	3.5%
Energy (\$0.30 per productive hour)	\$1,010	\$5,049	1.7%
Supervision and intervention (\$3.00 per productive hour)	\$10,098	\$50,490	17.5%
Total cost of ownership	\$57,818	\$289,089	100%

The \$75,000 purchase price is 25.9 per cent of the five-year cost of ownership. Supervision alone — the line most often omitted — is 17.5 per cent, two-thirds the size of the machine itself. Assumptions in Appendix A. Halving the purchase price to \$37,500 — with integration and maintenance, which are set as percentages of price, falling proportionally — reduces five-year cost of ownership by **31 per cent** to \$199,810 and cost per human-equivalent hour from \$24.54 to \$16.96. It does not halve either.

Exhibit 6 • Robot total-cost stack, dollars per productive hour, three configurations

Research question: where does the cost of an hour of machine labor actually sit? • **Units:** US dollars per productive robot hour, before the human-equivalence adjustment • **Period:** costed on 2025–26 input prices



Source: AUTHOR CALCULATION, Model 1. Wage inputs for supervision from BLS^{[9][10]}; utilization for Case C calibrated to BMW's published Figure 02 operating record^[25]. Full assumptions in Appendix A.


Takeaway: hardware capital is the largest single line only in the specialised industrial case, where integration is large but life is long and utilization high. In the general-purpose case, software, maintenance and supervision together exceed hardware. In the pilot-observed case every line inflates simultaneously because they are all divided by the same collapsed denominator — 1,088 productive hours a year instead of 3,366.

6. The robot cost curve

Robotics should get cheaper. The question is how fast, from which starting point, and in which cost line — because the industry's own recent price data does not show hardware deflation. Assigning batteries' or solar's learning rate to robots is the most common analytical error in this field, and this section explains why it fails.

Exhibit 7 • Comparative technology price declines, annualised over the stated period

Research question: do robots deflate like batteries, solar and semiconductors? • **Units:** annualised per cent change in price per year • **Period:** varies by series, stated on each row



Sources: BloombergNEF 2025 battery price survey^[24]; IFR quality-adjusted and nominal robot price series as compiled by Construction Physics^[44]; A3 North American robot order statistics for 2025 and H1 2026^{[42][43]}. Solar path is indicative. **Evidence status:** OBSERVED for battery packs, the 1990–2005 quality-adjusted robot series and the A3 order-value series; THIRD-PARTY ESTIMATE for the nominal robot series and the solar path. The A3 unit value is an AUTHOR CALCULATION: \$2.25bn / 36,766 units = \$61,197 in 2025; \$1.166bn / 17,995 units = \$64,796 in H1 2026.

Takeaway: robots did once deflate at battery-like rates — 83 per cent quality-adjusted between 1990 and 2005 — but that curve is thirty years old. The only current observed price series for the industry, average value per robot ordered in North America, rose 5.9 per cent between 2025 and the first half of 2026, because buyers are shifting toward higher-content systems. A deflation thesis for robotics must be argued on cost per unit of delivered work, not on unit price.

6.1 Which mechanisms actually apply

Mechanism	Applies?	Assessment
Manufacturing scale	Strong	Genuine, and largely unrealised: at ~19,100 humanoid units shipped worldwide in H1 2026 the category is at pre-Model-3 automotive volumes. ^[33] Tooling amortisation, yield and procurement leverage all improve steeply from this base. Agility's RoboFab is stated at up to 10,000 units a year of capacity ANNOUNCED . ^[29]
Component standardisation	Strong	Standard joint modules, connectors and interfaces cut both BOM and integration cost. Integration is 12 per cent of the base-case TCO, so this lever touches more than hardware.
EV supply chains	Partial	Shared motors, power electronics, thermal management and manufacturing know-how. But EV traction motors are optimised for continuous high-power output, not for the high-torque, low-speed, zero-backlash duty of a robot joint. The reducer, not the motor, is the constraint.
Batteries	Weak	Real deflation — \$108/kWh globally, \$84 in China in 2025 ^[24] — applied to a 1–3 kWh subsystem. A 13 per cent Chinese pack price fall saves roughly \$13–38 on a machine costing tens of thousands.
Sensors	Moderate	Cameras, depth and automotive LiDAR deflate on volumes robotics did not create. Offset by rising sensor counts: Figure 03 adds palm cameras and tactile sensing. ^[26]
Compute	Moderate	Performance per dollar improves fast; the absolute compute budget rises with model size. Jetson Thor's 7.5× compute over Orin arrives at 130 W and a \$3,499 developer-kit price. ^[41] Better inference buys task success, which is worth more than the saving.
Chinese manufacturing	Strong	The clearest observed price compression in the industry. Unitree reported 2025 revenue of ¥1.708bn with 62.9 per cent gross margin on humanoids COMPANY-REPORTED . ^[33] High margins at low prices imply low cost, and imply room to cut price further.
Learning curves	Unproven	No published quality-adjusted robot price series exists for the last decade. Claiming a learning rate requires a cumulative-volume price series; the industry does not publish one.
Software improvement	Strongest	The one lever that lowers cost per unit of work with no change to the BOM. In Model 5, moving uptime, task success and relative productivity from base to favourable cuts human-equivalent cost from \$24.54 to \$15.28 — a larger effect than cutting the purchase price from \$75,000 to \$30,000 AUTHOR CALCULATION .

WHY THE BATTERY AND SOLAR ANALOGIES MISLEAD

Solar modules and battery cells are **commodities sold by the watt and the kilowatt-hour**. Their learning curves are measured in a physical unit whose definition never changes, they are produced in enormous volume by a small number of process steps, and the buyer does not need to re-engineer a building to install them. Robots are **heterogeneous systems sold into heterogeneous workflows**. There is no denominator equivalent to the watt. Cost improvement in a robot can be spent on a better hand, more sensors or more compute instead of appearing as a lower price — and the observed A3 data suggests it currently is. The right learning-curve denominator for robotics is **delivered productive work**, not kilograms of hardware or nominal price per unit. On that denominator, cost is almost certainly falling; on the nominal-price denominator, the public data does not show it.

6.2 What a defensible cost-curve statement looks like

Three claims are supportable from current evidence. First, **Chinese platform prices are genuinely low and profitable**, which means the floor is lower than Western list prices suggest. Second, **the fastest cost decline available is in the denominator, not the numerator**: raising task success from 0.75 to 0.92 and relative productivity from 0.70 to 0.90 achieves more than any plausible near-term hardware saving. Third, **integration cost is the least-discussed and most tractable line**: it is 12 per cent of base-case TCO and up to 60 per cent of an industrial cell's installed capital, and it falls with standardisation and installer productivity rather than with component prices. Anything beyond these three claims is extrapolation.

7. Robot unit economics — Model 1

Model 1 converts a capital good into an hourly labor price. It is the spine of the report: Models 2, 3, 4 and 5 are all transformations of it. Its most important feature is the denominator, because that is where most public robot-cost claims fail.

MODEL 1 — ROBOT COST PER PRODUCTIVE HOUR

$$\text{CRF}(r,n) = r(1+r)^n / [(1+r)^n - 1]$$

$$C_R = [\text{CRF}(r,n) (P + I) + M + S + O] / (H \times U \times Q) + E + V$$

$$C_{R,HE} = C_R / \varphi$$

Term	Definition	Why it matters
<i>P</i>	Purchase price of the platform	The number everyone quotes. Ranks third of eleven in Model 5.
<i>I</i>	Integration investment	Guarding, fixtures, workflow change, commissioning. $0.2\text{--}2.0\times P$ depending on environment.
<i>r</i>	Financing rate	Ranks last of eleven. The cost of capital is not the constraint.
<i>n</i>	Useful life, years	Obsolescence-bounded, not wear-bounded, on fast-moving platforms.
<i>M</i>	Annual maintenance and repair	Set as a percentage of <i>P</i> ; scales with joint count and duty cycle.
<i>S</i>	Annual software, fleet and inference	Recurring; the vendor's ARR is the buyer's opex.
<i>O</i>	Annual insurance and other fixed	Rises with human proximity.
<i>H</i>	Scheduled hours per year	Hours the machine is rostered to work, not hours it exists.
<i>U</i>	Uptime = operating hours / scheduled hours	Charging, faults, blocked work, changeovers, waiting.
<i>Q</i>	Task-success / productive-time factor	Share of operating time that yields accepted output.
<i>E</i>	Energy cost per productive hour	Small per hour; a system constraint at fleet scale.
<i>V</i>	Allocated supervision per productive hour	Fully loaded supervisor cost divided by robots per supervisor.
φ	Productivity relative to the human benchmark	Converts a robot hour into a human-equivalent hour. Ranks first of eleven.

7.1 Three configurations, one formula


The report does not use a single base case, because "a robot" is not one thing. It uses three, each with a stated purpose. Case C is not a pessimistic scenario — it is the arithmetic that follows from the only multi-month operating record a manufacturer has published.

Input	A • Specialised industrial cell	B • General-purpose platform (mature)	C • Pilot-observed calibration	Basis
Purchase price <i>P</i>	\$75,000	\$75,000	\$75,000	Held constant deliberately, so the three cases differ only in how the machine is used
Integration <i>I</i>	150% of <i>P</i>	35% of <i>P</i>	35% of <i>P</i>	Industrial norm vs the general-purpose value proposition
Financing <i>r</i>	8%	10%	10%	Mature asset class vs emerging
Useful life <i>n</i>	8 years	5 years	4 years	Obsolescence risk rises with platform velocity
Maintenance <i>M</i>	6% of <i>P</i> /yr	12% of <i>P</i> /yr	15% of <i>P</i> /yr	Joint count, duty cycle, immaturity
Software <i>S</i>	\$3,000/yr	\$9,000/yr	\$12,000/yr	Fleet software and inference subscription
Insurance/other <i>O</i>	\$3,000/yr	\$2,000/yr	\$2,000/yr	Guarded cell carries more fixed compliance cost
Scheduled <i>H</i>	16h × 330d = 5,280h	16h × 330d = 5,280h	10h × 250d = 2,500h	Case C mirrors BMW's stated 10-hour weekday shift pattern ^[25]
Uptime <i>U</i>	0.92	0.85	0.58	Case C from ~1,250 operating hours ^[25] against ~2,150 scheduled
Task success <i>Q</i>	0.95	0.75	0.75	No vendor publishes this; a research assumption throughout
Energy <i>E</i>	\$0.60/h	\$0.30/h	\$0.30/h	~3 kW and ~1.5 kW average draw at \$0.20/kWh
Supervision <i>V</i>	\$1.80/h	\$3.00/h	\$6.00/h	\$45/h loaded technician at 1:25, 1:15 and 1:7.5 robots
Relative productivity φ	1.50	0.70	0.70	Purpose-built cell beats a human; general platform does not, yet
Productive hours/yr	4,615	3,366	1,088	$H \times U \times Q$
C_R	\$11.75	\$17.18	\$58.89	Per productive robot hour
$C_{R,HE}$	\$7.83	\$24.54	\$84.13	Per human-equivalent hour

AUTHOR CALCULATION throughout; inputs are **RESEARCH ASSUMPTIONS** except where sourced. Case C's *H* and *U* are calibrated to **COMPANY-REPORTED** BMW figures.^[25] Residual value is set to zero; taxes and task-specific facility costs are excluded unless captured in the integration percentage. Full assumption register in Appendix A.

Exhibit 8 • Model 1 output: robot cost per productive and per human-equivalent hour

Research question: what does an hour of machine labor cost, on comparable terms with an hour of human labor? • **Units:** US dollars per hour • **Period:** 2025–26 input prices



Source: [AUTHOR CALCULATION](#), Model 1. Human benchmark is fully loaded US warehouse labor, \$29.11/h, derived in Appendix C from BLS OEWS May 2025 and BLS ECEC March 2026. ^{[9][10]}

Takeaway: the same \$75,000 machine costs \$7.83, \$24.54 or \$84.13 per human-equivalent hour depending entirely on how it is used. The spread is 10.7× and none of it comes from the purchase price. This single result is the reason the report treats utilization and productivity, not hardware price, as the economically decisive variables.

WORKED DERIVATION, CASE B, SO THE ARITHMETIC CAN BE CHECKED

$CRF(0.10, 5) = 0.10 \times 1.10^5 / (1.10^5 - 1) = 0.161051 / 0.61051 = \mathbf{0.263797}$. $P + I = \$75,000 + \$26,250 = \$101,250$; annualised capital = $0.263797 \times \$101,250 = \mathbf{\$26,709}$. Annual fixed = $\$26,709 + \$9,000 (M) + \$9,000 (S) + \$2,000 (O) = \mathbf{\$46,709}$. Productive hours = $5,280 \times 0.85 \times 0.75 = \mathbf{3,366}$. $C_R = \$46,709 / 3,366 + \$0.30 + \$3.00 = \$13.88 + \$3.30 = \mathbf{\$17.18}$. $C_{R,HE} = \$17.18 / 0.70 = \mathbf{\$24.54}$.

7.2 The denominator is the discipline

Three errors recur in public robot-cost arithmetic, and all three live in the denominator.

Confusing scheduled with productive hours. A machine "capable of 24/7 operation" is quoted against 8,760 hours. The base case here yields 3,366 productive hours from 5,280 scheduled — and the pilot-observed case yields 1,088. Dividing the same annual cost by 8,760 instead of 3,366 makes a robot look 2.6× cheaper than it is.

Ignoring relative productivity. If two robots are needed to replace one worker's output, the robot labor requirement doubles and the cost per human-equivalent hour doubles with it. This is the ϕ term, and it is the largest single driver in the sensitivity analysis.

Setting supervision to zero. At a 1:15 robot-to-technician ratio, supervision is \$3.00 of the base case's \$17.18 per productive hour — 17.5 per cent of the total, and 17.5 per cent of five-year cost of ownership. Halve the ratio to 1:7.5 and C_R rises to \$20.18 and $C_{R,HE}$ to \$28.83, which erases the entire nominal advantage against \$29.11 US warehouse labor [AUTHOR CALCULATION](#). A vendor model that assumes autonomy has assumed away the cost that distinguishes a pilot from a business.

PART IV — ROBOTS VERSUS HUMAN LABOR**8. Robots versus human labor**

A wage is not a labor cost. In March 2026 US private-industry wages were **69.9 per cent** of employer compensation cost; the remaining 30.1 per cent is benefits and payroll burden, and even that excludes recruiting, turnover, absence cover and injury cost.^[9] Comparing a robot's fully loaded cost to a posted hourly wage overstates the robot's disadvantage by roughly 55 per cent. Comparing a robot's *purchase price* to a wage understates its cost by more.

8.1 The human cost stack

The BLS publishes the first two lines. The remaining three are author additions, deliberately conservative: 6 per cent of wage for recruiting and turnover, 4 per cent for absence backfill and overtime premium, and 2 per cent for injury and workers' compensation loading. The resulting loading factor is **1.551×** the posted wage.

Component	\$/hour at a \$18.77 median wage	% of fully loaded cost	Source / basis
Wages and salaries	\$18.77	64.5%	BLS OEWS May 2025 median, laborers and freight/stock/material movers ^[10] OBSERVED
Benefits and payroll burden	\$8.08	27.8%	BLS ECEC March 2026: wages are 69.9% of private-industry compensation ^[9] OBSERVED
Recruiting and turnover	\$1.13	3.9%	6% of wage RESEARCH ASSUMPTION
Absence backfill and overtime premium	\$0.75	2.6%	4% of wage RESEARCH ASSUMPTION
Injury, workers' compensation loading	\$0.38	1.3%	2% of wage RESEARCH ASSUMPTION
Fully loaded employer cost	\$29.11	100%	Loading factor 1.551× wage AUTHOR CALCULATION

Applied to other occupations in Appendix C. Shift constraints do not appear as a cost line: they appear as a cap on how many hours of a given wage a firm can buy at all, which is why Section 12 treats the operating window separately.

8.2 Symmetry: what each side actually carries

Dimension	Human labor	Machine labor
Cost structure	Almost entirely variable. Stops when the work stops.	Almost entirely fixed. Continues whether or not the work exists — which is why demand variability is a robot's worst enemy and a temp worker's normal condition.
Capacity flexibility	Hired and released on weeks; overtime absorbs peaks at a premium.	Fixed for the life of the asset. A robot bought for a peak is idle in a trough and its cost per productive hour rises accordingly.
Operating window	Bounded by shift law, fatigue, willingness and premium pay.	Bounded by charging, maintenance and upstream supply — not by fatigue. This is the genuine structural advantage, and Section 12 quantifies how much of it is real.
Task range	Enormous. A warehouse associate does twenty things without re-training.	Narrow, and re-scoping costs engineering time. Vision-language-action models are attacking exactly this constraint.
Failure mode	Distributed. One person's bad day costs one person's output.	Concentrated. One controller fault stops a line; downtime is correlated across identical units.
Quality	Variable within and between individuals; improves with tenure.	Highly repeatable within the trained distribution; degrades sharply outside it, often silently.
Safety and liability	Injury cost, workers' compensation, regulatory exposure.	Product liability, functional-safety certification, insurance rising with human proximity.
Data	Produces no structured process record by default.	Produces a complete, timestamped process record. This is a genuine and undervalued second product.
Risk that does not exist on the other side	Turnover, absenteeism, hiring market tightness, industrial action.	Obsolescence, vendor failure, cyber exposure, spare-parts supply, single-supplier component risk.

8.3 Comparing quality-adjusted output, not nominal hours

The comparison that matters holds accepted output constant. Two mechanisms break the nominal-hour comparison in opposite directions.

Against the robot: if a machine produces 70 per cent of a trained human's accepted output on the same task, then 1.43 robot hours are needed per human hour, and a \$17.18 robot hour becomes a \$24.54 human-equivalent hour. This is the ϕ term. Against a \$29.11 benchmark it clears on nominal arithmetic and fails a 30 per cent switching margin, which is the honest reading of the general-purpose case today.

For the robot: in a purpose-built cell the same logic runs the other way. At $\phi = 1.5$ an \$11.75 robot hour becomes a \$7.83 human-equivalent hour — a 3.7× advantage over the same \$29.11 benchmark. Quality consistency, scrap reduction and process data are additional and are not counted here at all, which makes the industrial case conservative.

FOUR ASYMMETRIES THAT NO HOURLY-COST COMPARISON CAPTURES

Irreversibility. A hire can be reversed in weeks; a \$101,250 installed capital commitment cannot. Firms therefore require a margin of safety, not equality — this report uses 30 per cent, consistent with a payback requirement of roughly two years on the incremental investment.

Capacity, not cost, is often the buying reason. If a firm cannot hire at any wage, the relevant comparison is not robot cost versus wage but robot cost versus foregone output. Section 14 treats this case, which is economically different and politically much more comfortable.

Option value. A deployed fleet gets cheaper per unit of work through software updates with no change to the BOM. A hired workforce does not appreciate in the same way, though it does accumulate tacit knowledge that the robot has to be taught explicitly.

The human is the fallback. Every current robot deployment retains people to handle exceptions. The cost of that fallback is the V term, and it is what turns a technically working robot into an unprofitable one.

9. Human-robot cost parity — Model 2

Parity is not the point at which robot cost equals wage. It is the point at which the robot's cost per human-equivalent hour is enough below the employer's fully loaded marginal labor cost to justify an irreversible capital commitment. This report applies a 30 per cent switching margin throughout.

MODEL 2 — THE BREAK-EVEN PURCHASE PRICE

Adoption condition: $C_{R,HE} < W/k$ where W = fully loaded human cost, $k = 1.30$

Setting $I = iP$ and $M = mP$ and solving for price:

$$P^* = [(\varphi W - E - V) \times (H U Q) - (S + O)] / [CRF(r, n)(1 + i) + m]$$

9.1 The parity matrix

Exhibit 9 • Human-robot parity matrix: purchase price against fully loaded wage

Research question: at what combination of hardware price and local labor cost does machine labor actually win? • **Units:** ratio of fully loaded human cost to robot cost per human-equivalent hour • **Period:** 2025–26 input prices

Ratio of human cost to robot human-equivalent cost — values above 1.30 clear a 30% switching margin

		\$10	\$15	\$20	\$25	\$35	\$50	\$75
robot purchase price (resulting $C_{R,HE}$)	\$150k (\$39.69/h)	0.25	0.38	0.50	0.63	0.88	1.26	1.89
	\$100k (\$29.59/h)	0.34	0.51	0.68	0.84	1.18	1.69	2.53
	\$75k (\$24.54/h)	0.41	0.61	0.81	1.02	1.43	2.04	3.06
	\$50k (\$19.49/h)	0.51	0.77	1.03	1.28	1.80	2.57	3.85
	\$30k (\$15.45/h)	0.65	0.97	1.29	1.62	2.27	3.24	4.85
	\$20k (\$13.42/h)	0.75	1.12	1.49	1.86	2.61	3.73	5.59
		\$10	\$15	\$20	\$25	\$35	\$50	\$75
fully loaded human cost per hour								

Source: **AUTHOR CALCULATION**, Model 2 on the Case B chassis: 16h × 330d scheduled, U 0.85, Q 0.75, φ 0.70, n 5, r 10%, integration 35%, maintenance 12% of price, S \$9,000, O \$2,000, E \$0.30, V \$3.00.

Takeaway: at the \$29.11 US warehouse benchmark a \$75,000 general-purpose platform sits at a ratio of about 1.19 — nominally cheaper, but inside the switching margin, which is a decision to wait rather than a decision to buy. The same platform is decisively economic at \$35/h and above, and decisively uneconomic below \$20/h at any price above \$50,000. Cheap-labor geographies are not the humanoid market, and that is an arithmetic result, not a preference.

9.2 Break-even purchase price, by wage and operating regime

The same equation solved for price makes the trade-off explicit. Each cell is the highest purchase price at which the platform exactly meets the fully loaded wage in that column, before the switching margin. Negative values mean no price works: even a free machine loses money because recurring costs alone exceed the labor it displaces.

Operating regime	\$10/h	\$15/h	\$20/h	\$25/h	\$35/h	\$50/h	\$75/h
1 shift (8h), U 0.85, Q 0.75, ϕ 0.70	−\$10,024	\$2,347	\$14,719	\$27,091	\$51,834	\$88,949	\$150,808
2 shifts (16h), U 0.85, Q 0.75, ϕ 0.70	\$3,054	\$27,798	\$52,541	\$77,284	\$126,771	\$201,002	\$324,719
2 shifts (16h), U 0.92, Q 0.90, ϕ 0.85	\$24,644	\$63,668	\$102,692	\$141,716	\$219,763	\$336,835	\$531,955
3 shifts (20h), U 0.92, Q 0.90, ϕ 0.85	\$36,581	\$85,360	\$134,140	\$182,920	\$280,480	\$426,820	\$670,719

Exhibit 9b. Research question: how much machine can a given wage environment justify? Units: US dollars of purchase price. Period: 2025–26 input prices. Source:

AUTHOR CALCULATION, Model 2; other inputs as Case B. Values are break-even before the 30% switching margin; applying the margin reduces every cell by roughly 40–50 per cent. **Takeaway:** operating regime rivals the wage bracket. Doubling the shift window at a \$25 wage raises the affordable price from \$27,091 to \$77,284, close to what doubling the wage itself from \$25 to \$50 achieves on a single shift (\$88,949) — and a firm controls its shift pattern in a way it does not control the local wage. Adding software gains in uptime, task success and relative productivity is worth a further \$64,432 at the same \$25 wage.

9.3 What the sensitivity actually shows

Financing matters least. Sweeping the discount rate from 3 to 18 per cent moves human-equivalent cost by \$4.35 per hour, the smallest range of eleven variables tested. Subsidised finance is not what unlocks robotics.

Useful life matters, but the assumption is the risk. Extending life from five to eight years cuts the capital recovery factor from 0.2638 to 0.1874 — annualised capital falls 29 per cent. The risk is that a platform generation on a two-year cadence makes an eight-year assumption a write-down rather than a saving. Case A can justify eight years because a welding cell's task does not change; a general-purpose platform cannot yet.

Wage level is necessary but not sufficient. Germany, Japan and Korea have both high labor cost and the highest robot densities, consistent with the model. But the United States has a similar wage level and one-quarter of Korea's density,^[2] which the wage variable alone cannot explain — industry mix, plant scale, integrator supply and management capability all enter through I , U and Q rather than through W .

WHERE PARITY ALREADY EXISTS, ON THIS REPORT'S OWN ARITHMETIC

Cleared, with margin. High-throughput repetitive tasks in engineered environments — automotive body-in-white, electronics handling, machine tending, palletising, structured inspection, tote transport in a mapped building. Case A prices these at \$7.83 per human-equivalent hour against \$29–39 fully loaded human cost. This is not a forecast; it is the reason 4.66m industrial robots exist.

At the boundary. General-purpose platforms in high-wage two-shift logistics and manufacturing: \$24.54 against \$29.11, clearing nominally and failing the margin. This is precisely where the pilots are, and it is why they are pilots.

Not cleared, and not close. Open-ended construction, bedside care, household work, irregular repair, delicate harvesting. Here either the duty cycle collapses the denominator (Section 10) or the capability gate has not been passed at any price.

10. The \$10-per-hour robot — Model 3

This section is scenario analysis, not forecast. It asks a conditional question: if a general-purpose platform reached \$20, \$10, \$5, \$2 or \$1 per human-equivalent hour at reference utilization, which tasks would enter the automation frontier? The answer is non-obvious, because the reference cost is not the cost a task actually pays.

10.1 Why a reference price is not a task price

A platform quoted at L dollars per human-equivalent hour achieves that figure only when it can be worked at the reference duty cycle. Real tasks cannot supply continuous work: a restaurant has two peaks a day, a strawberry field has a few weeks a year, a construction site is accessible when the other trades are not there. And real tasks impose an exception burden — human intervention, rework and re-teaching — that scales the cost up again.

MODEL 3 — THE AUTOMATION FRONTIER

$$C_{task} = L \times (1 / \lambda) \times x$$

Task enters the frontier $\Leftrightarrow C_{task} < W_{FL} / 1.30$ **and** capability gate ≤ 2


Equivalently, the break-even reference price for a task is $L^* = (W_{FL} / 1.30) \times \lambda / x$

Symbol	Meaning	Range used and why
λ	Duty-cycle attainment: the share of the reference operating window for which the task can actually supply work	1.00 for continuous line work; 0.50 for large-area cleaning and construction transport; 0.35 for hospitality peaks; 0.25 for field spraying; 0.15 for delicate harvest; 0.10 for household tasks
x	Exception burden multiplier: uplift for intervention, rework, re-teaching and recovery	1.05 for structured inspection; 1.15–1.25 for transport and packing; 1.50 for each-picking and trailer unloading; 1.80–2.00 for detail cleaning, occupied-space trades and household tidying
Capability gate	1 passed today; 2 near, paid pilots exist; 3 uncertain, demos only; 4 blocked primarily by non-cost factors	Assigned from the deployment evidence ladder in Section 4.1, not from vendor claims

All three parameters are **AUTHOR FRAMEWORK** and are set out task by task in Appendix B. They are calibrated to the deployment record, not fitted to a result.

Exhibit 10 • The automation frontier: break-even reference platform cost by task

Research question: how cheap must machine labor be before each task becomes economic, and does cheapness alone suffice? • **Units:** US dollars per human-equivalent hour at reference utilization (log scale) • **Period:** wages May 2025; readiness August 2026




Source: **AUTHOR CALCULATION** and **AUTHOR FRAMEWORK**, Model 3. Fully loaded human costs derived from BLS OEWS May 2025 medians and the BLS ECEC compensation split. ^{[9][10]} Capability gates assigned from the evidence in Section 4.1.

Takeaway: the tasks span a 40-fold range in break-even platform cost, from \$26.32 for spot welding to \$0.66 for household tidying — and the ordering is driven by duty cycle and exception burden, not by wage. Every green dot (capability already proven) to the right of \$10 is a market that exists today. Every red dot is a task where a cheaper robot changes nothing.

Exhibit 10b • Frontier width at five reference cost levels

Research question: how many tasks enter as cost falls, and what stops the count from rising? • **Units:** count of tasks from a 25-task screen • **Period:** as above



Source: **AUTHOR CALCULATION**, Model 3, on the 25-task screen in Appendix B. Scenario only; not a forecast.

Takeaway: the economic gate opens steadily — 2, 10, 18, 22, 23 tasks at \$20, \$10, \$5, \$2 and \$1 — but the joint gate saturates at 18. Below \$5 per hour, cost stops being the constraint. Five of the 25 tasks are economically viable at a dollar an hour and remain undeployable because of dexterity, liability, regulation or trust. This is the report's central quantitative finding about the limits of cheap machine labor.

10.2 Reading the frontier by sector

Sector	\$20	\$10	\$5	\$2–\$1	What the numbers say
Manufacturing	2 of 5	4	4	4	Already the frontier. Spot welding breaks even at \$26.32 and structured inspection at \$20.94; dexterous mixed-model assembly breaks even at \$12.37 but sits behind capability gate 3, so it never enters however cheap the platform.
Warehousing	0 of 4	3	4	4	Tote movement enters at \$15.26, packing at \$12.34, each-picking at \$11.17, trailer unloading at \$7.57. Trailer unloading is late not because it is hard but because trailers arrive in bursts — $\lambda = 0.55$.
Agriculture	0 of 3	0	1	2	Milking enters at \$7.68 because cows are milked on a schedule. Field spraying needs \$3.09 and delicate harvest \$1.34, purely because of seasonality. Agriculture is a duty-cycle problem wearing a robotics costume.
Construction	0 of 3	0	2	2	Material transport at \$7.68 and rebar tying at \$5.74 enter around \$5. Finish trades break even at \$5.46 but sit behind gate 4 — they never enter on cost.
Hospitality & retail	0 of 4	1	3	4	Shelf scanning enters at \$10.79; back-room replenishment at \$7.61; food running needs \$4.25. The 11 per cent decline in hospitality service-robot sales in 2024 is what a λ of 0.35 looks like in a revenue line. ^[3]
Cleaning	0 of 2	0	1	1	Large-area scrubbing enters at \$8.00 — and professional cleaning grew 34 per cent in 2024, the fastest service segment. ^[3] Detail cleaning needs \$4.09 and is behind gate 3.
Healthcare	0 of 2	1	1	1	Supply transport enters at \$12.95. Patient handling breaks even at \$4.16 and is gate 4: the highest-injury task in the sector will be among the last automated, and cost is not why.
Home services	0 of 2	0	0	0	Laundry folding needs \$0.88 and tidying \$0.66. At $\lambda = 0.10$, a \$10 platform costs \$150–200 per useful hour in a house. The household robot is not gated on intelligence; it is gated on the fact that a house has almost no work for it.
Total of 25 tasks	2	9	16	18	The joint gate saturates at 18; seven tasks never enter at any price.

Exhibit 10c. Counts are of tasks passing *both* gates at each level. Source: **AUTHOR CALCULATION**, Model 3; task screen in Appendix B. Illustrative scenario — not a forecast of prices or of adoption.

THREE THINGS THE FRONTIER DOES NOT SAY

It is not a forecast that platforms will reach these prices. Section 6 shows the only current observed price series moving the other way. The frontier is conditional arithmetic.

It is not a claim about jobs. Every row is a task. A warehouse associate performs several of these rows plus a dozen that are not screened here. Automating four of a worker's twenty tasks changes the job before it removes it — see Section 19.

It does not model the demand response. If machine labor at \$5 an hour made a service dramatically cheaper, demand for it would rise, which raises λ and improves the economics further. That feedback is real, is second-order relative to the capability gate, and is deliberately excluded to keep the model legible.

PART V — PRODUCTIVITY ECONOMICS

11. Robotics and productivity

The empirical literature on robots and productivity does not agree, and the disagreement is not noise. Different papers use different units of observation, different exposure measures and different equilibrium assumptions, and they answer different questions. Presented fairly, they constrain the thesis rather than confirming or refuting it.

11.1 The contradiction matrix

Study	Unit & period	Headline finding	What it can and cannot support
Graetz & Michaels (2018) <i>Rev. Econ. Stat.</i> [13]	17 countries, 14 industries, 1993–2007	Robot densification raised annual labor productivity growth by −0.36pp — about 15% of aggregate productivity growth. TFP rose about two-thirds as much. Output prices fell significantly. No significant effect on aggregate hours; hours of low-skilled workers fell relative to others. Wages captured ~10% of the productivity gain.	Supports: robots raise measured productivity and lower output prices. Cannot support: claims about local labor markets, or about the post-2007 period. Industry-country panel; no within-country geography.
Acemoglu & Restrepo (2020) <i>J. Polit. Econ.</i> [12]	US commuting zones, 1990–2007	One more robot per thousand workers reduced the employment-to-population ratio by 0.18–0.34pp and wages by 0.25–0.50% . One robot in a local labor market coincided with an employment drop of 5.6 workers .	Supports: concentrated local displacement. Cannot support: an aggregate national employment conclusion — the design deliberately nets out national general-equilibrium offsets, which is the point of the identification and also its limit.
Dauth, Findeisen, Suedekum & Woessner (2021) <i>J. Eur. Econ. Assoc.</i> [14]	German regions and workers, 1994–2014	One additional robot per 1,000 workers displaced about 1.7 manufacturing jobs, fully offset by roughly 2.0 service jobs . No significant regional wage effect. Incumbent workers' plant tenure rose; young labor-market entrants bore the adjustment.	Supports: reallocation rather than net destruction, under strong labor-market institutions. Cannot support: transfer to a US-style institutional setting — German works councils and internal redeployment are part of the mechanism.
Aghion, Antonin, Bunel & Jaravel (2020) France [15]	19,448 plants, 245 industries, 1994–2015	Firm-level employment elasticity to automation of +0.2 on impact rising to +0.4 after ten years; positive for high-, medium- and low-skill alike. Industry-level employment elasticity +0.41 , but +0.40 only where international competition is high and +0.02 (insignificant) where it is low. Producer prices −0.05 ; profits +0.78 .	Supports: automation as a competitiveness technology that expands the automating firm. Cannot support: a general claim — the effect is conditional on trade exposure, which is the paper's own most important result and the one most often dropped when it is cited.
OECD (2019, 2022, 2024) [16][18][51]	Cross-country, occupational and regional	Automation lowers employment in elementary and some middle-skill occupations and raises it in professional and technical roles; regions with high automation risk often recorded productivity gains without aggregate employment decline; firms adopt cobots when local hiring constraints bind.	Supports: compositional change with aggregate stability, and a scarcity-driven adoption motive. Cannot support: inference about any individual worker or place — these are averages that conceal concentrated loss.
NY Fed / BLS productivity data [11]	US manufacturing, 1987–2022	Manufacturing labor productivity growth fell from +3.4%/yr (1987–2007) to −0.5%/yr (2010–2022) ; TFP from +1.4% to +0.1% . The slowdown is pervasive across industries and firm sizes. Fast-growing "leader" industries fell from +6.5% to −0.6% .	Supports: the strongest single counterexample to robot-driven productivity optimism. Cannot support: a causal claim that robots caused the slowdown — it is a puzzle the authors explicitly leave open.

Exhibit 13. Research question: what does the empirical record actually establish about robots, productivity, employment and wages? Units: as stated per row. Period: 1987–2022 across studies. Evidence status: **OBSERVED DATA** and peer-reviewed estimates. **Takeaway:** the findings are reconcilable but not combinable. Productivity gains, aggregate employment stability and concentrated local wage losses can all be true at once, because they are measured at different levels of aggregation.

11.2 How the contradictory findings fit together

The three-effect decomposition explains the apparent conflict without discarding any result. A robot **displaces** tasks in the adopting plant; the resulting cost reduction raises output and demand for complementary labor, the **productivity effect**; and new tasks appear in integration, maintenance, data work and supervision, the **reinstatement effect**. The displacement effect is local, fast and visible. The productivity and reinstatement effects are diffuse, slow and often land somewhere else — in another city, another industry, or in a lower price paid by a consumer who never learns why.

A commuting-zone design measures the first effect cleanly and the other two hardly at all, by construction. A cross-country industry panel measures the net of all three and therefore finds productivity gains with stable hours. A firm-level design in a trade-exposed industry finds the productivity effect dominating, because the automating firm takes share from a non-automating competitor — which is a transfer,

not a national gain, and is exactly why the French result vanishes where international competition is low.

11.3 The measured-productivity problem

Between 2015 and 2024 the world's operational robot stock nearly tripled, from 1.63m to 4.66m, and US installations ran at 34,000–40,000 units a year.^{[1][5][7]} Across 2010–2022, almost the same window, US manufacturing labor productivity growth was *negative*.^[11] Four explanations are consistent with the evidence, and they are not mutually exclusive:

- **Composition.** Robots concentrate in automotive and electronics; the measured slowdown was steepest in exactly those "leader" industries, which suggests the driver is something else — plausibly the exhaustion of the semiconductor quality-adjustment effect in the price deflators.
- **Mismeasurement.** Quality improvement, variety and reliability gains delivered by automation are imperfectly captured in output deflators, so real output is understated.
- **Scale.** 4.66m robots is large in absolute terms and small relative to a global manufacturing workforce of hundreds of millions. Graetz and Michaels' 0.36pp was a real but modest contribution; a contribution of that size can be swamped by other forces.
- **Complementary investment.** Robots deliver productivity only alongside process redesign, management capability and worker skill. Where they are bolted onto an unchanged process, they raise capital intensity without raising output per hour.


The ILO's cross-country work on robots, employment and trade reaches a similarly qualified conclusion: the effects run through trade and reallocation channels rather than through a simple domestic substitution of machines for workers.^[49] The honest conclusion is that **robot adoption is neither a sufficient nor a reliably observable cause of aggregate productivity growth**. It is a mechanism that works at the plant level and gets lost in the aggregate. Any thesis claiming that scaling machine labor will visibly lift national productivity statistics must explain the last fourteen years, and none of the sources reviewed here does so convincingly.

12. The 24/7 economy — Model 4

A robot's ability to work outside human shift constraints is its clearest structural advantage. It is also the one most often overstated, because extending the scheduled window degrades the uptime and task-success factors that convert it into productive hours. Model 4 quantifies the difference between the theoretical and the realistic version of that advantage.

Exhibit 11 • Utilization economics: scheduled, operating and productive hours

Research question: how much of the 24/7 advantage survives realistic uptime, charging, maintenance and intervention? • **Units:** US dollars per human-equivalent hour (left); hours per year (right) • **Period:** 2025–26 input prices, 330 operating days



Source: AUTHOR CALCULATION, Model 4 on the Case B chassis with annual fixed cost held at \$46,709 in all four regimes.

Takeaway: moving from one shift to a theoretical 24-hour day cuts human-equivalent cost from \$39.82 to \$21.43 — a 46 per cent reduction, not the 58 per cent the naive calculation gives. Over four-fifths of the entire benefit (\$15.28 of \$18.39) is captured by the first step, from 8 to 16 hours. Beyond 20 hours the curve is nearly flat, because charging, maintenance and changeovers have to happen somewhere and the degradation of U and Q offsets the extra scheduled time.

Regime	Scheduled h/yr	Uptime U	Operating h/yr	Success Q	Productive h/yr	Fixed cost per scheduled hour	C_R	$C_{R,HE}$
1 shift, 8h/day	2,640	0.90	2,376	0.80	1,901	\$17.69	\$27.87	\$39.82
2 shifts, 16h/day	5,280	0.85	4,488	0.75	3,366	\$8.85	\$17.18	\$24.54
3 shifts with maintenance window, 20h/day	6,600	0.80	5,280	0.72	3,802	\$7.08	\$15.59	\$22.27
24h theoretical	7,920	0.72	5,702	0.70	3,992	\$5.90	\$15.00	\$21.43
Naive: U and Q both held at 0.90	2,640→7,920	0.90	—	0.90	2,138→6,415	—	—	\$35.92→\$15.12

Exhibit 11b. **AUTHOR CALCULATION**, Model 4. The declining U and Q across regimes are **RESEARCH ASSUMPTIONS** reflecting that longer windows have to absorb charging, preventive maintenance and unstaffed periods; they are the model's most contestable inputs and are swept in Model 5.

12.1 The three distinct benefits of a longer window

Fixed cost per scheduled hour falls, mechanically. Annual fixed cost of \$46,709 divided by 2,640 scheduled hours is \$17.69; divided by 7,920 it is \$5.90 — a 67 per cent reduction. This is the part of the argument that is arithmetically guaranteed.

Cost per productive hour falls, but by less. The same move takes \$27.87 to \$15.00, a 46 per cent reduction, because productive hours grow only from 1,901 to 3,992 — $2.1\times$ against a $3.0\times$ increase in the scheduled window. The gap between 67 and 46 per cent is the entire practical content of Model 4.

Return on the surrounding plant rises, and this is usually the larger prize. A robot's own capital is small next to the building, the machine tools, the racking and the utility connection it uses. If a \$20m plant currently runs one shift, extending to two roughly halves its fixed cost per unit of output regardless of what the robots themselves cost. In capital-intensive processes this effect dominates the robot's own economics, and it is the strongest argument for automation in semiconductors, metals, chemicals and high-value machining.

12.2 Where the bottleneck migrates

Bottleneck	Why it binds when the window extends, and what it costs
Upstream supply	A cell starves if inbound material only arrives on a day shift. Extending robot hours without extending logistics converts uptime into waiting, which shows up as a lower U , not as a cost line.
Downstream capacity	Packaging, quality release, palletising or shipping must also run. A robot blocked at the outfeed is operating and unproductive — the Q term.
Maintenance windows	Preventive maintenance is not optional; at 24 scheduled hours it has to be taken out of production, which is why U falls to 0.72 in the fourth regime.
Charging and energy	A mobile platform charges. Agility states 16 hours of operation per charge for Digit v4 and about 22 for v5 COMPANY-REPORTED ^[29] — which caps a 24-hour window unless the fleet is oversized or hot-swapped, and oversizing the fleet is another way of raising cost per productive hour.
Supervision coverage	Night-shift technician coverage is thinner and more expensive. The V term rises exactly when the window that was supposed to lower cost is extended.
Demand	The decisive one. Running 24 hours produces 24 hours of output. If the market absorbs 16, the extra capacity becomes inventory, and inventory is a cost. Utilization economics only work where demand is genuinely unconstrained or where the plant is capacity-limited.
Energy price and grid	Off-peak power is usually cheaper, which helps. At fleet scale the constraint becomes interconnection rather than tariff — see Section 20.

THE OBSERVED CASE, AND WHY IT MATTERS

BMW's Figure 02 pilot ran **ten-hour weekday shifts** — the most conservative regime in the table — and achieved an operating ratio near 0.58 within it. ^[25] This is not a criticism of the deployment; a pilot inside a running body shop is correctly scheduled around production rather than around the robot. But it means the industry's flagship humanoid operating record was generated under the *least* favourable utilization regime the model contains, and the economics it implies (\$84.13 per human-equivalent hour) are therefore the economics of a pilot, not of a business. The path from \$84 to \$24 runs almost entirely through this section: more scheduled hours, higher uptime within them, and a supervision ratio that improves as the deployment matures.


PART VI — DEMOGRAPHICS AND LABOR SCARCITY

13. Demographics and labor scarcity

Demographic decline raises the shadow price of reliable labor. It does not automatically create robot demand. The OECD projects old-age dependency ratios reaching unprecedented levels across member countries over the next 35 years and warns that GDP per capita growth will slow significantly in most of them without changes in policy and behaviour.^[17] But the chain from a shrinking working-age population to an installed robot has five links, and each of them can break.

Exhibit 14 • Working-age population (15–64), indexed to 2000

Research question: which economies face labor scarcity from the supply side, and how far has it already gone? • **Units:** index, 2000 = 100 • **Period:** 1990–2024 observed



Source: World Bank World Development Indicators, series SP.POP.1564.TO, retrieved 29 August 2026.^[21] **Evidence status:** OBSERVED DATA — these are estimates of what has already happened, not projections. Forward-looking statements below are UN WPP 2024 medium-variant projections and are labelled as such.^[20]

Takeaway: this is not a future problem for four of the six economies. Japan's working-age population peaked in 1995 and is 16.6 per cent below that peak; Italy peaked in 2011, China in 2015, Germany in 2016 and Korea in 2019. Only the United States has not peaked, and its continued growth is driven by immigration. The countries with the highest robot densities are, with the exception of Singapore, precisely the countries furthest past their working-age peak.

Economy	Peak year	Peak (m)	2024 (m)	Change from peak	Robot density 2024	UN WPP 2024 medium-variant outlook ^[20]
Japan	1995	87.5	72.9	–16.6%	446	Total population projected to fall by 21 million, 2024–2054 — the second-largest absolute loss in the world
Italy	2011	39.2	37.4	–4.6%	Top 20	Among 19 countries where net immigration has delayed or reduced the population peak
China	2015	988.4	976.8	–1.2%	166	Largest absolute loss in the world: –204 million by 2054; –786 million by 2100
Germany	2016	54.0	52.5	–2.7%	449	Immigration is delaying the peak; without it the decline is steeper
Korea	c. 2019	37.5	36.3	–3.2%	1,220	Among the world's lowest fertility rates; steep projected decline
United States	not yet	—	220.1	still rising	307	One of 52 countries where immigration is the main driver of population growth through 2054

OBSERVED DATA for peaks and 2024 values (World Bank WDI^[21]) and for robot density (IFR^[2]). Outlook column is the UN's medium-variant projection^[20] and is a projection, not an observation. Note that China's working-age decline to date is small; its projected total population decline is the largest in the world. These are different statements and are often conflated.

13.1 The five-link chain, and where it breaks

#	Link	Evidence, and the condition required for the link to hold
1	Working-age population falls	OBSERVED in Japan, Italy, China, Germany and Korea. Holds .
2	→ labor becomes scarce in the specific occupation	Conditional . A falling aggregate population does not imply scarcity in warehousing if participation, migration, retirement age or sector reallocation adjust. Japan's response has included raising participation among older workers and women as much as it has included robots.
3	→ wages and vacancies rise	Conditional . Scarcity shows up as wage pressure only where the employer cannot substitute across locations. Traded manufacturing can move the plant; a hospital cannot move the patients.
4	→ the automation calculation flips	Follows Model 2 mechanically — a rise from a \$25 to a \$35 fully loaded wage raises the affordable purchase price from \$77,284 to \$126,771 on the two-shift regime AUTHOR CALCULATION .
5	→ robots are actually installed	The weakest link . Requires finance, an available integrator, a task the technology can do, management capability and an acceptable payback. Japan's installations <i>fell</i> for a fourth consecutive year in 2024 to 44,500, and Germany's fell 5 per cent ^[1] — both while their working-age populations continued to shrink.

THE COUNTEREXAMPLE THAT DISCIPLINES THE DEMOGRAPHIC ARGUMENT

Japan has the longest-running working-age decline of any large economy, one of the highest robot densities in the world, and a domestic robot industry among the best in the world. Its robot installations declined in 2024 for the **fourth consecutive year**.^[1] Germany, past its working-age peak since 2016, installed 5 per cent fewer robots in 2024 than in 2023. Meanwhile India — young, growing and low-wage — grew installations 7 per cent, and China, whose working-age population has barely fallen, took 54 per cent of the world's new robots.

Demography is a real and durable force on the demand for automation. But over a five-year horizon, the capital cycle, the industrial mix, the state of the customer's end market and the availability of integrators explain robot installations far better than the age structure does. Anyone forecasting robot demand from demographic projections should reconcile that forecast with these four observations first.

14. Robots as labor-scarcity technology

The framing matters economically and politically. A labor-*replacement* technology substitutes for a worker who would otherwise be employed. A labor-*scarcity* technology fills work that would otherwise not be done at all. A labor-*complementing* technology raises the output of a worker who stays. The same machine can be all three in different settings, and which one it is determines both the economics and the politics of its adoption.

Setting	What the robot substitutes for	Primary character	Evidence
Manufacturing	Hard-to-staff night shifts, heavy or hazardous handling, and increasingly a marginal worker who is available	Mixed — scarcity in high-wage ageing economies, replacement in mature high-density plants	OECD case evidence documents manufacturers adopting cobots when local hiring constraints became binding ^[48] ; Acemoglu & Restrepo document displacement in US commuting zones ^[12] . Both are true, in different places.
Agriculture	Seasonal crews that are increasingly unavailable at any legal wage	Scarcity	The strawberry system in Section 4.2 targets a stated \$43,000 per acre per year of picking labor and one operator supervising eight machines — a complement, not a replacement, at that ratio ^[50]
Logistics	Peak-season surge labor that is hired and released within a quarter	Scarcity at peak, replacement at base	Amazon hired 250,000 seasonal workers for the 2025 holiday while introducing new robotic systems ^[28] ; the same firm's internal planning reportedly targets avoiding 600,000 future US hires ^[48] . The distinction between not-hiring and firing is the whole political argument.
Healthcare	Non-clinical transport and lifting; not clinical judgment	Complementing	Intuitive's installed base grew 12% to 11,106 systems in 2025 with procedures up 17% in Q4 — the surgeon is not replaced, the procedure count rises ^[47]
Elderly care	Lifting, transport, monitoring — not companionship or judgment	Scarcity, in principle	The clearest theoretical case for scarcity-driven robotics: Japan, Italy, Korea. Yet Section 10 scores patient handling as gate 4 — blocked by liability, consent and trust. Demand exists and supply does not, which is why this remains the largest unserved robotics market in the world.
Construction	Material handling on sites where skilled trades are scarce	Scarcity	Trade shortages are acute in most advanced economies, but Section 10 puts material transport's break-even at \$7.68 — the duty cycle, not the shortage, is binding.

14.1 Why the distinction changes the analysis

It changes the counterfactual. If work is not done because no one can be hired, the robot's value is the output it creates, which is usually far larger than the wage it displaces. This is why scarcity-driven deployments can clear a payback test that replacement-driven deployments fail.

It changes the political economy. "We could not fill the shift" and "we cut the shift" produce identical robot counts and completely different labor-market consequences. Not-hiring is invisible in displace-

ment statistics, which is why Acemoglu–Restrepo-style designs may understate the long-run employment effect of automation in a tightening labor market and Dauth-style designs may overstate the ease of reallocation.

It changes what to measure. The right indicators for a scarcity-driven adoption are unfilled vacancy duration, output foregone, injury rates and wage levels for the workers who remain — not headcount. A policy debate conducted only in jobs saved or lost will misread most of what is currently happening.

PART VII — THE TWO SYSTEMS

15. China's robotics advantage

China's advantage is not a single thing. It is the coincidence of the world's largest deployment market, the deepest component supply chain, and near-monopoly control of the one input every actuated joint needs. Separating what has been observed from what has been announced is essential here, because the announcements are unusually loud.

15.1 What has actually been observed

- **Deployment.** 295,000 industrial robots installed in 2024, **54 per cent** of the world total; operational stock above **2 million**; the largest market for twelve consecutive years **OBSERVED**. ^{[1][46]}
- **Domestic share.** Chinese suppliers took **57 per cent** of Chinese installations in 2024, up from 47 per cent in 2023 — a ten-point shift in one year, and in metal and machinery applications the local share is very high **OBSERVED**. ^{[1][46]}
- **Electronics concentration.** China hosted **64 per cent** of global electronics-sector robot installations in 2024 **OBSERVED**. ^[46]
- **Production.** China produced **595,000 industrial robots** and 13.5 million service robots in the first nine months of 2025, exceeding its full-year 2024 output **OBSERVED**, reported by the China Machinery Industry Federation. ^[45]
- **Humanoid shipments.** Chinese manufacturers accounted for **97 per cent** of the ~19,100 humanoid units shipped globally in H1 2026 **THIRD-PARTY ESTIMATE**. ^[33]
- **Profitable low prices.** Unitree reported 2025 revenue of ¥1.708bn (~\$253m) with an overall gross margin of **60.3 per cent** and **62.9 per cent** on humanoids, and adjusted net profit near ¥600m, ahead of a Shanghai STAR listing at a ~¥61bn market capitalisation **COMPANY-REPORTED**. ^[33] The G1 lists from **\$13,500**. ^[32]
- **Critical materials.** China holds **60 per cent** of magnet rare-earth mining, **91 per cent** of refining and **94 per cent** of sintered permanent magnet production, up from about 50 per cent in 2005 **OBSERVED**. ^[22]

A DEFINITIONAL TRAP: PRODUCTION IS NOT INSTALLATION

China's reported **595,000 industrial robots produced** in nine months of 2025^[45] exceeds the IFR's count of **542,076 installed worldwide** in all of 2024.^[1] Both numbers are real. They are not comparable. The Chinese statistic counts manufactured units on a national statistical definition that includes simpler robotic arms and manipulators; the IFR counts installations of machines meeting ISO 8373 on a survey of suppliers and users. Treating the two as the same series produces the false conclusion that China alone installs more robots than the world does. Registered as Contradiction C2.

15.2 Why the component position is the durable advantage

The subsystem that dominates a robot's bill of materials — actuators, motors and precision reducers — depends on high-energy permanent magnets, and permanent magnets depend on separated heavy rare earths. The IEA's assessment is that outside China, existing and announced capacity meets only about 50 per cent of projected 2035 mining demand, 25 per cent of refining and **under 20 per cent** of magnet manufacturing, and that roughly **\$60 billion** of investment through 2035 would be needed to diversify.^[22] The IEA further estimates that full implementation of export controls would expose about **\$6.5 trillion** of annual downstream production value outside China.

The control regime has been active and volatile: initial controls on seven heavy rare earths in April 2025, a major expansion in October 2025, a one-year suspension in November 2025, and tightened

dual-use controls affecting Japan in January 2026.^[22] For a robotics programme, this is not a background geopolitical risk. It is a direct input to the largest line of the bill of materials, and it is the single most concentrated dependency in the entire value chain.

15.3 What has been announced but not observed

Chinese national and provincial programmes have set humanoid and embodied-AI targets that this report does not treat as output. Third-party shipment forecasts for the category — roughly 60,000 units globally in 2026 and 500,000 by 2030 THIRD-PARTY ESTIMATE^[33] — are forecasts. The gap between shipment and deployment is the point: the same source estimates roughly **250 Unitree units in active deployment** against more than 5,500 shipped in 2025. China's demonstrated near-term advantage is cost and volume in *specialised and semi-general* systems, not a proven general-purpose humanoid workforce.

16. America's Physical AI advantage

The United States leads in the layers above the machine: foundation models, accelerators, cloud, simulation, developer tooling, venture finance and universities. The economic question is whether that leadership converts into robot labor — and the answer depends on whether frontier capability or rapid diffusion of sufficient capability is what matters.

16.1 The American stack, and the one deployment asset

- **Foundation models for control.** Gemini Robotics 1.5 and Gemini Robotics-ER 1.5 (Google DeepMind), $\pi_{0.5}$ (Physical Intelligence), and the Open X-Embodiment collaboration are all US-led or US-anchored.^{[36][38][40]}
- **Edge compute.** NVIDIA's Jetson AGX Thor is the de facto reference platform, with Agility, Amazon Robotics, Boston Dynamics, Caterpillar, Figure, Hexagon, Medtronic and Meta named as early adopters and 1X, John Deere, OpenAI and Physical Intelligence evaluating; NVIDIA cites over 2 million developers on its robotics stack COMPANY-REPORTED.^[41]
- **Simulation and synthetic data.** The cheapest available scaling lever, and predominantly American software.
- **Capital.** Agility's \$620m SPAC transaction at a \$2.5bn pre-money valuation, with Foxconn leading a ~\$200m PIPE, is representative of the depth of American risk capital for pre-revenue robotics COMPANY-REPORTED.^[29]
- **The one large operational fleet.** Amazon's one millionth robot, deployed July 2025 across 300+ facilities, with the DeepFleet foundation model claimed to improve fleet travel time by 10 per cent COMPANY-REPORTED.^[27] This is the decisive counterexample to the "America does software, China does hardware" framing.

16.2 Frontier capability or rapid diffusion?

This is the strategically decisive question and the evidence points toward diffusion. Three arguments.

First, capability is not the binding constraint in the tasks that are actually economic. Every task in the frontier screen that clears both gates at \$10 per hour is a structured task that a purpose-built system already performs. Frontier reasoning helps at the margin; it does not decide whether tote transport is profitable.

Second, models diffuse faster than hardware does. Gemini Robotics-ER 1.5 shipped through a public API in September 2025.^[36] A capability that reaches an API reaches every integrator with a subscription, including those buying Chinese hardware. Frontier leadership in a diffusible layer converts to durable advantage only where the hardware, the deployment channel or the data is also controlled.

Third, the scarce American resource is deployment, not intelligence. The United States installed 34,200 industrial robots in 2024, down 9 per cent, at a density of 307 against Korea's 1,220.^{[1][2]} The gap is not a model gap. It is a gap in integrators, in plant scale, in capital cycle and in management practice — the terms I , U and Q in Model 1, not the terms Physical AI improves fastest.

Exhibit 15 • China–United States robotics ecosystem comparison

Research question: which layers of the robotics value chain does each system control, and how contestable is each? • **Units:** qualitative position with the supporting observed figure • **Period:** evidence to August 2026

Layer	China	United States	Contestability
Frontier models for control	Fast follower; strong applied research	Leads — Gemini Robotics 1.5, $\pi_{0.5}$, Open X-Embodiment	Low barrier to diffusion once shipped through an API. Advantage decays quickly.
Accelerators and edge compute	Constrained by export controls; domestic substitution under way	Leads decisively — Jetson Thor is the reference platform	Policy-dependent both ways. The most politically exposed layer.
Cloud and simulation	Large domestic capacity	Leads	Moderate. Simulation quality is a research race, not a capital race.
Robot manufacturing	Leads — 595k units produced in 9M 2025; 97% of H1 2026 humanoid shipments	Rebuilding; Agility's RoboFab stated at up to 10,000 units/yr capacity ANNOUNCED	Low. Manufacturing scale takes years and is where China's lead compounds.
Actuators, reducers, encoders	Rapid localisation; cost leadership	Thin; largely imported from Japan and China	Low. The most durable Chinese advantage after magnets.
Permanent magnets & rare earths	Dominant — 94% of sintered magnets, 91% of refining	Minimal; ex-China capacity meets <20% of 2035 magnet demand	Very low over five years; ~\$60bn and a decade to diversify.
Deployment market	Leads — 54% of global installations	34,200 installations in 2024, ~9%	Low. Demand follows the manufacturing base.
Large operating fleets & fleet data	Large but less publicly documented	Leads — Amazon, 1m+ robots, 300+ facilities	Low, and proprietary. The most defensible American position.
Venture capital	State-directed and abundant	Leads — \$620m SPAC at \$2.5bn pre-money for a pre-scale humanoid firm	High. Capital is the most mobile input of all.
Integrators and installers	Deep, because volume created them	Thin and ageing	Low in the short run. Trained people are the slowest input to add.


Sources: IFR ^{[1][2][46]}; China Machinery Industry Federation ^[45]; IEA ^[22]; NVIDIA ^[41]; Amazon ^[27]; Agility ^[29]; humanoid shipment estimates ^[33]. **Evidence status:** figures as tagged; the "leads" judgments and contestability column are **AUTHOR FRAMEWORK**. **Takeaway:** the layers America leads are the most diffusible; the layers China leads are the least. That asymmetry, not the current model leaderboard, is what determines who captures the economics of machine labor.

PART VIII — THE FLYWHEEL**17. The robotics flywheel — a link-by-link failure audit**

The flywheel argument is that deployment generates data, data improves policies, better policies raise task success, higher return supports volume, volume lowers cost, and lower cost expands deployment. Each arrow is a conditional statement. This section states the condition, then tests it against evidence. Six of the nine links are currently unproven.

Exhibit 16 • The robotics flywheel and its failure points

Research question: is the deployment→data→capability→cost→deployment loop self-reinforcing? • **Units:** qualitative causal diagram • **Period:** assessed August 2026



Source: **AUTHOR FRAMEWORK**, with the supporting evidence for each link in the table below.

Takeaway: the loop is plausible and no link is impossible. But a nine-link loop in which each link holds with probability 0.8 turns with probability 0.13. Flywheel arguments should be assessed on the weakest link, not on the strength of the story, and the weakest links here are data portability, capability-to-ROI conversion, and cost falling into price rather than into product content.

#	Link	Status	Condition required, and the evidence for and against
1	Deployment → real-world data	Holds	Deployed robots do generate data. Amazon's fleet trains DeepFleet on warehouse movement ^[27] ; Agility reports 65,000+ cumulative operating hours ^[29] . Uncontroversial.
2	Data → better training	Partial	Requires transferability. Open X-Embodiment demonstrates real cross-embodiment gains — RT-1-X +50%, RT-2-X ~3× on emergent skills ^[38] — but across 22 heterogeneous embodiments the transfer is far from universal. $\pi_{0.5}$ scales in the number of <i>environments</i> , and environments are collected by people in buildings, not scraped. ^[40] Data from one site and one body does not automatically improve another.
3	Training → higher capability	Holds	Well documented on benchmarks: RT-2 raised unseen-scenario performance from 32% to 62% ^[37] ; Gemini Robotics-ER 1.5 leads 15 spatial benchmarks and transfers motion across embodiments ^[36] .
4	Capability → broader task coverage	Partial	Requires that capability be the binding constraint. Model 3 shows it usually is not: below \$5 per hour, five of 25 screened tasks are cost-viable and capability-blocked, but many more are blocked by duty cycle, which no model improves. Capability opens some tasks and is irrelevant to others.
5	Task coverage → higher ROI	Unproven	Requires that added tasks raise utilization more than they raise exception cost. A robot that does five tasks poorly can have a lower Q than one that does one task well. No public deployment has published a before/after intervention rate. This is the single largest evidence gap in the whole industry.
6	ROI → greater demand	Partial	Requires ROI to clear a margin, not parity. Agility reports \$300m+ of multi-year pre-order commitments and 30+ pipeline customers COMPANY-REPORTED ^[29] . Against that, Amazon — the most capable operator in the world — withdrew Blue Jay from operations in February 2026 ^[28] , and A3 North American unit orders grew only 2.0% in H1 2026 ^[42] .
7	Demand → production scale	Partial	Requires supply chain to follow. Chinese humanoid output is scaling fast ^{[33][45]} , but actuators, precision reducers and rare-earth magnets are concentrated bottlenecks: ex-China magnet capacity meets under 20% of projected 2035 demand ^[22] . Volume can be gated by a component, not by orders.
8	Scale → lower cost	Unproven	Requires cost reduction to reach the price. Unitree's 62.9% humanoid gross margin ^[33] shows there is room to cut price. The A3 series shows average value per robot ordered rising 5.9% into H1 2026 ^{[42][43]} . Manufacturers currently spend cost improvement on capability, not on price.
9	Lower cost → more deployment	Partial	Requires cost to be the binding constraint. Model 3's central result is that below about \$5 per hour it is not. Nine tasks enter the frontier between \$20 and \$10; only two more enter between \$5 and \$1.

Exhibit 16b. Status column is **AUTHOR FRAMEWORK**: *holds* = supported by multiple independent sources; *partial* = supported under stated conditions that do not always obtain; *unproven* = no public evidence either way at the level of paid production.

17.1 The two competing routes to fleet learning

Route	Example	Strength and weakness
Proprietary homogeneous scale	Amazon: one million robots of a small number of designs, in buildings the operator controls ^[27]	Strength: enormous, consistent, labelled operational data with a single owner who captures the whole gain. Weakness: generalises poorly outside that building type, and the owner has no commercial reason to share it.
Collaborative pooling	heterogeneous Open X-Embodiment: 1m+ trajectories, 22 embodiments, 60 datasets, 34 labs ^[38]	Strength: breadth across bodies and tasks; demonstrated positive transfer. Weakness: research-grade, heterogeneous, and thin on the failure and recovery cases that matter commercially.

Neither route has produced evidence of the smooth power-law scaling associated with language models, and there is a structural reason to doubt it will. Language-model scaling laws hold because tokens are fungible and the corpus is effectively free. Robot trajectories are neither. The realistic expectation is **step-function improvements within task families** — a manipulation class becomes reliable, then another — rather than a smooth curve across all physical work.

17.2 The teleoperation trap

Teleoperation makes a deployment work before autonomy does, and supplies exactly the rare recovery data that autonomy needs. It is therefore both the bridge and the trap. Every teleoperated minute is a variable human labor cost inside a product sold as capital equipment — the V term in Model 1. A vendor that scales deployments faster than it reduces intervention minutes scales its own cost of goods with revenue and never reaches the operating leverage the model promises. The diligence question that follows is gate 6 of the Investment Opportunity Map, and it is the one that separates a robotics business from a staffing business with unusually expensive staff.

PART IX — MACROECONOMICS

18. The deflationary machine

Robotics can lower unit production cost. Whether that reaches a consumer price depends on four things it does not control: how much of unit cost is labor, how much of that labor is addressable, how much of the addressable labor cost actually falls, and how much of the saving competition forces into price rather than margin. Multiply four fractions and the result is small.

18.1 The transmission chain


COST-TO-PRICE TRANSMISSION

$$\Delta P_{consumer} = \sigma_L \times \alpha \times \delta \times \rho$$

labor share of unit cost × addressable share of that labor × cost reduction achieved on it × pass-through to price

Exhibit 17 • Cost-to-consumer-price transmission by sector

Research question: if machine labor halves the cost of the labor it can address, how much reaches a price? • **Units:** per cent • **Period:** illustrative scenario, no date-specific claim



Source: **AUTHOR FRAMEWORK** and **AUTHOR CALCULATION**. Labour shares, addressable shares and pass-through rates are **RESEARCH ASSUMPTIONS** chosen to show the structure of the chain, with a uniform 50 per cent assumed reduction in the cost of the addressable labour. They are not estimates of any specific market.

Takeaway: even on generous assumptions the largest net effect in the screen is a **5.2 per cent** price reduction in parcel and contract logistics, and the smallest is **0.3 per cent** in hospital care. Robotics is a plausible source of a one-off level shift in the price of a few goods and services. It is not a general deflationary force, because in most of the economy labour is a minority of unit cost and the automatable part of that labour is a minority again.

18.2 Where cost reduction is most likely to reach a price

Traded manufactured goods. Competitive, elastic supply, low labor share — and the one sector where the effect is empirically documented. Graetz and Michaels find robot densification reduced output prices significantly, and Aghion and co-authors find an industry-level producer-price elasticity of **-0.05**, statistically real but small.^{[13][15]} Note that the French result also shows a profit elasticity of **+0.78**: most of the gain stayed with firms.

Logistics. The highest labor share in the screen at roughly 42 per cent, with a large addressable component and reasonably competitive markets. If robotics compresses prices anywhere visibly, it will be here first. Morgan Stanley's estimate that Amazon could save **\$2-4 billion a year from 2027**, and Amazon's own reported internal figure of up to **\$10 billion**, are cost savings, not price reductions **THIRD-PARTY ESTIMATE**.^[48] Whether any of it reaches a delivery price depends on competition in last-mile delivery, not on the robots.

Food processing. Moderate labor share, high addressability, but final food prices are dominated by agriculture, energy, packaging, logistics and retail margin. The robot addresses a slice of a slice.

18.3 Where it will not

Construction. Land, zoning, permitting, materials, finance and regulation dominate the delivered price of a building. Even complete automation of on-site material transport — which Model 3 puts at a \$7.68 break-even — addresses a small share of total project cost.

Healthcare. Labor share is the highest in the screen at roughly 55 per cent, and the addressable share is the lowest. Clinical labor is licensed, and reimbursement, not cost, sets price in most systems. Intuitive's experience is instructive in the opposite direction: robotic surgery expanded procedure volumes 17 per cent year on year and grew revenue to \$10.06bn — it raised throughput and quality, and there is no evidence it lowered the price of surgery.^[47]

Anything land- or licence-constrained. Where the binding input is a permit, a plot or a professional monopoly, cheaper labor raises the rent on the constrained input rather than lowering the price.

FOUR REASONS COST DECLINE DOES NOT BECOME PRICE DECLINE

Market concentration. A concentrated market retains the saving as margin. Aghion's +0.78 profit elasticity against -0.05 on prices is what that looks like in data.^[15]

Quality absorption. Firms often spend the saving on a better product rather than a lower price — the same behaviour Section 6 documents inside robots themselves, where falling component costs appear as more sensors rather than lower prices.

Demand elasticity. If cheaper output expands demand, added capacity is absorbed at an unchanged price. This raises welfare and output without showing up in the CPI.

Constrained complements. Energy, land, commodities, permits and regulated inputs do not automate. As machine labor scales, their relative price rises — which is the constraint-migration argument of Section 20 and the reason a machine-labor economy is not automatically a low-price economy.

19. Labor, wages and inequality

The Acemoglu–Restrepo decomposition — displacement, productivity, reinstatement — is the right framework because it explains why the empirical literature disagrees without requiring any of it to be wrong. The three effects operate at different speeds, at different places, and on different people.

Effect	Mechanism	Empirical anchor	Speed, visibility and incidence
Displacement	A task previously performed by a person is performed by capital; labor demand for that task falls.	US commuting zones: -0.18 to -0.34 pp on employment-to-population and -0.25 to -0.50% on wages per robot per thousand workers; ~ 5.6 workers per robot locally ^[12]	Fast, local, highly visible. Falls on specific workers in specific plants in specific towns. Politically salient precisely because it is legible.
Productivity	Lower unit cost raises output, lowers prices, raises real incomes and raises demand for complementary labor.	+0.36pp of annual labor productivity growth, $\sim 15\%$ of aggregate; output prices fall; wages capture $\sim 10\%$ of the gain ^[13]	Slow, diffuse, nearly invisible. Lands on consumers who never attribute it, and on workers in other industries.
Reinstatement	New tasks appear that did not exist: integration, fleet operations, maintenance, data work, safety engineering, exception handling.	German evidence: ~ 2.0 service jobs created per ~ 1.7 manufacturing jobs displaced ^[14] . Amazon states 700,000+ employees upskilled since 2019 and 30% more technical roles at its newest fulfillment centre ^[27] COMPANY-REPORTED	Slowest and most institution-dependent. Requires training, mobility and time. Does not reliably reach the displaced worker.

Exhibit 18. Research question: how do the three effects of automation differ in speed, visibility and incidence? Units: as stated per row. Period: 1990–2026 across sources.

Takeaway: the net employment effect is an empirical result, not a mechanical consequence. But the *distribution* is not symmetric: the displacement effect is concentrated and immediate while the offsetting effects are dispersed and delayed, which is why the aggregate can be neutral while the experience is not.

19.1 Jobs versus tasks — the distinction that changes the number

Every row of the Section 10 frontier is a task. A warehouse associate performs several of them plus many that are not screened. Automating four of a worker's twenty tasks does three things simultaneously: it removes some labor demand, it raises the productivity of the remaining sixteen, and it changes what the job is. Occupation-level automation forecasts routinely conflate these, which is how a 50 per cent "task automatability" score becomes a headline about half of all jobs disappearing.

The IMF's exposure framework handles this correctly by splitting exposure from complementarity: about **60 per cent** of employment in advanced economies is exposed to AI, but that splits into roughly **27 per cent** high-exposure/high-complementarity and **33 per cent** high-exposure/low-complementarity; in emerging markets exposure is 40 per cent and in low-income countries 26 per cent.^[19] Exposure is not displacement. Complementarity is where the sign of the effect is decided.

19.2 Distribution: wages, labor share and who captures the gain

Wages capture little of the productivity gain directly. Graetz and Michaels put wage gains at roughly 10 per cent of the productivity improvement.^[13] Aghion and co-authors find no significant firm-level average wage effect and no significant change in the labor share, with a profit elasticity of +0.78.^[15] Both point the same way: the direct gain accrues mainly to capital, with the indirect gain reaching workers as consumers through lower prices — which, per Section 18, is a weaker channel than it sounds.

Ownership concentration determines the split. If machine labor is owned by a few large operators with proprietary fleets, the productivity effect accrues to those owners and to high-skill complements. If

robots-as-a-service lowers the entry barrier for small and medium manufacturers — and IFR reports service-robot RaaS growing 31 per cent in 2024^[3] — ownership broadens and more of the gain reaches downstream firms and consumers. Business model, not technology, decides this.

The incidence falls on entrants. The German evidence is the most specific: incumbent workers' plant tenure *rose* under robot exposure, while young labor-market entrants absorbed the adjustment through reduced manufacturing hiring.^[14] Automation in an institutional setting with internal redeployment protects those inside and closes the door to those outside. That is a different distributional problem from mass displacement, and it needs different policy.

19.3 What would change this assessment

Three observable developments would justify revising the labor-market conclusions upward in severity: a published intervention rate below one per cent per task hour in a paid multi-customer deployment; a general-purpose platform demonstrating ϕ above 1.0 across five or more task families in one facility; or a sustained decline in average value per robot ordered in the A3 series alongside rising unit volumes. None of the three has been observed as of the research cut-off.

PART X — JUDGMENT

20. Conclusion: what becomes scarce?

The central question was whether robotics can reduce the marginal cost of physical labor the way AI is reducing the marginal cost of certain forms of intelligence. The answer is that it can, for a defined and expanding subset of physical work, and that the mechanism is structurally different from the software case in a way that will not go away.

20.1 Why intelligence and physical labor scale differently


A trained model has a large fixed cost and a marginal cost of replication close to zero. Every additional unit of machine labor requires an additional machine: manufactured from constrained materials, powered, transported, integrated into one specific workflow, maintained, insured and supervised. Section 5 shows the purchase price is a quarter of five-year cost of ownership, and the other three quarters are irreducibly physical and local.

This is why the report's answer is conditional rather than binary. The intelligence layer of robotics genuinely does behave like software: a model improvement propagates across a fleet at near-zero marginal cost and, per Model 5, improving U , Q and ϕ together is worth more than halving hardware price. But the intelligence layer sits on top of a physical layer that obeys manufacturing economics, and the composite scales at the speed of the slower layer.

20.2 If machine labor scales, what binds next?

Exhibit 19 • Constraint migration map

Research question: which constraints bind today, and which would bind if machine labor scaled tenfold? • **Units:** qualitative score, 2 = slack to 5 = binding • **Period:** today (2026) versus a 10× scenario



Source: **AUTHOR FRAMEWORK**, informed by IEA rare-earth and electricity assessments^{[22][23]}, IFR density and installation data^{[1][2]} and the deployment evidence in Section 4. Scores are qualitative and are a scenario, not a forecast.

Takeaway: the two constraints that bind hardest today — wage parity and task capability — are the two that loosen most in the scenario. Everything that gets worse is upstream or adjacent: magnets, actuators, integrators, grid interconnection, standards and capital. A machine-labor economy does not remove scarcity; it relocates it from the labor market to the supply chain and the electricity system.

Constraint	Why it binds at scale	The best available quantification
Electricity and grid	Robots and the inference that drives them both need power and both need an interconnection, which takes years.	Global electricity demand forecast to grow 3.6% a year to 2030 , at least 2.5× the pace of total energy demand; US data centres alone drive around half of American demand growth; grid investment must rise roughly 50% from \$400bn a year ^[23] OBSERVED / IEA FORECAST
Permanent magnets and rare earths	Every actuated joint needs a high-energy magnet; there is no substitute at comparable torque density.	China: 94% of sintered magnets, 91% of refining. Ex-China capacity meets <20% of projected 2035 magnet demand; ~\$60bn of investment needed ^[22] OBSERVED
Actuators, reducers, encoders	The largest BOM block, a concentrated supply base, and precision manufacturing that does not scale on software timescales.	No public capacity series exists — itself a finding, and one of the report's registered evidence gaps
Integrators and installers	Integration is 12–60% of installed capital and is delivered by trained people who take years to produce.	US density 307 vs Korea 1,220 at comparable wages ^[2] is consistent with a deployment-capacity constraint rather than a cost constraint OBSERVED
Safety, standards, liability	Every model update to a machine that shares space with people is a safety-relevant change requiring validation.	Five of 25 screened tasks are cost-viable at \$1/hour and blocked on non-cost grounds AUTHOR CALCULATION
Edge compute and semiconductors	Each robot carries its own inference; fleet-scale deployment is fleet-scale silicon demand.	Jetson AGX Thor at 130 W and \$3,499 for a developer kit ^[41] OBSERVED
Capital	Machine labor converts an operating expense into a capital commitment; scaling it means financing a very large asset base.	Financing rate ranks last of eleven variables in Model 5 — the constraint is availability of balance sheet, not the cost of capital
Downstream demand	Capacity without demand is inventory. Section 12's utilization gains require an end market that absorbs the output.	Hospitality service-robot sales fell 11% in 2024 ^[3] — a demand-side, not a capability-side, result OBSERVED
Real-world data	Loosens over time as fleets grow — the one constraint the flywheel genuinely relieves.	$\pi_{0.5}$: performance rose monotonically from 3 to 104 training environments ^[40] OBSERVED

20.3 Where human comparative advantage moves

The frontier screen answers this directly, and the answer is not "creativity". The tasks that resist automation longest in the screen are not the most cognitively demanding: they are household tidying, patient transfer, detail cleaning, finish trades and delicate harvest. What they share is **unstructured environments, low duty cycle, high consequence of failure, and a requirement for consent or trust**. Human comparative advantage moves toward work with those properties, toward judgment under accountability, toward relationships, and toward defining new tasks — the reinstatement effect of Section 19.

20.4 Final judgment

VERDICT: PARTIALLY SUPPORTED

The thesis — *robotics may transform selected forms of physical labor from scarce human inputs into scalable capital services* — is **supported for structured, high-duty-cycle, low-consequence physical tasks** and **not supported for general physical work**.

What is established. Machine labor at \$7.83 per human-equivalent hour against \$29–39 of fully loaded human cost, in a category with 4.66 million operating units. That is scalable capital service, and it already exists.

What is plausible but unproven. That general-purpose platforms reach \$10 per human-equivalent hour. Model 5 shows the path requires four levers to move at once; no single lever suffices, and the only current observed price series moves the wrong way.

What is not supported. That cheap machine labor generalises to all physical work. Below \$5 per hour the binding constraint changes character: five of 25 screened tasks stay closed at a dollar an hour. That is a statement about capability, liability and duty cycle, and no price solves it.


What would change this verdict. A published intervention rate below 1 per cent per task hour in a paid multi-customer deployment; a demonstrated ϕ above 1.0 for a general-purpose platform across several task families; or a sustained fall in average value per robot ordered alongside rising volumes. The first would move the verdict to *supported*. Their continued absence three years from now would move it toward *rejected*.

Model 5 • Robot cost sensitivity matrix

Model 5 sweeps eleven inputs across defensible ranges, one at a time, and then in combination, to identify which variables actually determine the cost of machine labor. Base case is Case B at \$24.54 per human-equivalent hour.

Exhibit 12 • Sensitivity of human-equivalent cost to each input

Research question: which variables move the cost of machine labor, and by how much? • **Units:** US dollars per human-equivalent hour • **Period:** 2025–26 input prices



Source: AUTHOR CALCULATION, Model 5. Each bar spans the outcome across that variable's full tested range with all others at base; ranges are listed in Appendix A.


Takeaway: the ranking is the finding. Relative productivity (\$30.67 of range), scheduled hours (\$26.43) and purchase price (\$26.27) lead; financing rate (\$4.35) is last. Three of the top five — productivity, uptime and task success — are software-and-operations variables that improve without changing the bill of materials.

Rank	Variable	Range tested	Best	Worst	Range	% of base
1	Relative productivity ϕ	0.40 – 1.40	\$12.27	\$42.94	\$30.67	125.0%
2	Scheduled hours/day H	8 – 24 h	\$17.93	\$44.36	\$26.43	107.7%
3	Purchase price P	\$20,000 – \$150,000	\$13.42	\$39.69	\$26.27	107.1%
4	Uptime U	0.50 – 0.97	\$22.09	\$38.42	\$16.33	66.5%
5	Task success Q	0.50 – 0.97	\$20.04	\$34.45	\$14.41	58.7%
6	Integration, % of P	10% – 150%	\$22.44	\$34.19	\$11.75	47.9%
7	Supervision V	\$0.25 – \$8.00 /h	\$20.61	\$31.68	\$11.07	45.1%
8	Useful life n	3 – 12 years	\$19.51	\$30.48	\$10.97	44.7%
9	Software S	\$1,000 – \$24,000 /yr	\$21.14	\$30.90	\$9.76	39.8%
10	Maintenance, % of P /yr	3% – 25%	\$21.67	\$28.68	\$7.01	28.6%
11	Financing rate r	3% – 18%	\$22.59	\$26.94	\$4.35	17.7%

Exhibit 12b. AUTHOR CALCULATION. One-at-a-time sweeps around the Case B base of \$24.54.

Exhibit 12c • What it takes to reach \$10 and \$5 per human-equivalent hour

Research question: can any single lever deliver a \$10 robot hour? • **Units:** US dollars per human-equivalent hour • **Period:** illustrative scenario



Source: **AUTHOR CALCULATION**, Model 5. Levers: hardware = P to \$30,000; software = U 0.92, Q 0.92, ϕ 0.90; utilization = 22 scheduled hours/day; supervision = V \$0.75/h. The final two bars add an 8-year life, 6% maintenance and \$3,000 software, then a \$20,000 price with 15% integration and 5% financing. **Illustrative scenario — not a forecast.**

Takeaway: no single lever reaches \$10. Hardware alone gives \$15.45, software alone \$15.28, utilization alone \$19.13 and supervision alone \$21.32. Two levers reach \$9.95; three reach \$8.24; four reach \$5.74. The \$10 robot is a compound outcome requiring simultaneous progress on price, reliability, scheduling and autonomy — and each of those has an independent probability of not happening.

MARKET OUTLOOK

Robotics Market Outlook and Valuation

How large is the robotics industry today, how large could it become, and how much of that depends on Physical AI? This section sizes the market bottom-up from observed unit counts and observed price anchors, then projects it to 2030 with a longer view to 2035 and 2040 under three scenarios. It values an *industry*, not a company: the question is how big a business machine labor can become, not what any share of it is worth.

WHY THIS IS BUILT BOTTOM-UP RATHER THAN TAKEN FROM A MARKET-SIZE STUDY

Published robotics market sizes range from roughly \$50bn to well over \$200bn for the same year, because they measure different things: the robot alone, the robot plus its cell, the robot plus software and services, or the whole automation capital-goods complex. The IFR — the only body that counts robots unit by unit — puts the **global industrial robot installations market at an all-time high of US\$16.7bn** on 542,076 units, which is an average of **\$30,807 per robot at the robot level** OBSERVED DATA ^{[1][52]}. Any figure several times larger is counting something else. This section therefore builds the market from units and prices it can name, states each layer separately, and lets the reader see exactly which definition produces which number.

M.1 What the market is today: a bottom-up 2025 baseline

Segment	2025 US\$bn	Share	How it is built, and from what
Industrial robots robot level only	17.5	16.3%	IFR's \$16.7bn on 542,076 units in 2024, grown to the IFR's 575,000-unit 2025 outlook at a \$30,400 average OBSERVED ^{[1][52]}
Integration, peripherals and cells	35.0	32.6%	Installed system cost runs 2–3× the robot itself; ^[44] a 2.0× multiplier is applied, the conservative end of that range ASSUMPTION
Service, spares and maintenance	16.5	15.4%	A ~5.9m-unit professional installed base at roughly \$2,800 a year of contracted service, spares and field labour — about 6% of installed system value ASSUMPTION
Medical and surgical robots	14.5	13.5%	Intuitive alone reported \$10,064.7m of 2025 revenue; ^[47] the balance is Stryker, Medtronic, Johnson & Johnson, Globus, CMR and rehabilitation and laboratory systems OBSERVED + ASSUMPTION
Consumer service robots	8.2	7.6%	Close to 20 million units sold in 2024, overwhelmingly floor cleaners and mowers, ^[3] at a blended retail price near \$375
Robotics software, AI and simulation	5.3	4.9%	Fleet management, simulation, perception and foundation-model services sold separately from hardware, at roughly \$900 per professional unit a year ASSUMPTION
AMR and logistics robots	5.3	4.9%	102,900 transport and logistics service robots sold in 2024 ^[3] at a \$45,000 blended price, plus fleet controls
Other professional service robots	4.3	4.0%	42,000 hospitality, 25,000 professional cleaning, 19,500 agriculture and 3,100 inspection units ^[3] at segment-typical prices
Humanoid robots	0.6	0.6%	~12,000 units shipped in 2025 at a blended \$50,000 — a price range spanning Unitree's \$13,500 list ^[32] to Agility's \$125,000 disclosed bill of materials ^[29] 3RD-PARTY EST.
Total global robotics market, 2025	107.2	100%	AUTHOR CALCULATION from the components above

Collaborative robots are a subset of industrial robots and are not counted separately — they were 19.6% of North American units ordered in 2025 at a \$33,417 average order value. ^[43] Warehouse conveyance, sortation and storage systems are excluded: they are automation capital goods, not robots on the IFR definition, and including them would roughly double the total. **Two observations follow immediately.** First, the robots themselves are the *smallest* part of the robotics market — robot hardware of all kinds is 47 per cent of the total, and industrial robots alone are 16 per cent. Second, at \$107bn the entire global robotics industry is about one per cent of what the world spends on the manufacturing and logistics labour it competes with, which is the arithmetic reason this market can grow for a very long time before it saturates.

M.2 The growth model: how each factor enters the arithmetic

Projections are not a CAGR extended forward. Each segment is either **unit-driven** (annual units × average selling price) or **base-driven** (cumulative installed base × annual spend per unit), and the report's structural findings enter through named variables. The chain runs as follows.

THE MARKET-EXPANSION CHAIN

technology improvement → robot capability → falling cost per task → higher adoption → larger installed base → new revenue pools → market expansion


$$\begin{aligned} \text{Unit-driven: Revenue}_t &= \text{Units}_t \times \text{ASP}_t \quad \cdot \quad \text{Base-driven: Revenue}_t = \text{InstalledBase}_t \times \text{SpendPerUnit}_t \\ \text{InstalledBase}_t &= \text{InstalledBase}_{t-1} \times (1 - 0.05) + \text{Units}_t \end{aligned}$$

Structural factor	Where it enters the model	Mechanism, and the evidence behind the assumption
Robot deployment and installed-base growth	Annual unit paths for every hardware segment; the installed-base roll-forward at a 5% retirement rate	The engine of the whole model. The professional installed base grows from 5.9m to 10.7m units by 2030 in the base case — and because service and software revenue are base-driven, every unit added keeps paying for a decade. ^[1]
Falling cost per robot	ASP paths: industrial \$30.4k→\$24k, AMR \$45k→\$22k, humanoid \$50k→\$13k by 2040	Falling price is a headwind to market value and a tailwind to market size. Section 6 shows the only current observed price series rising, so ASP decline is modelled as gradual for mature segments and steep only for humanoids, where volume is 37× below scale. ^{[42][43]}
Physical AI and foundation models	Software spend per professional unit: \$900 → \$2,200 (2030) → \$3,700 (2040)	The single largest new revenue pool. Model 5 showed uptime, task success and relative productivity dominate cost per productive hour — so the software that moves them can charge for a share of the value it creates, per robot per year. ^{[36][38][40]}
Robots performing more tasks	Unit paths in AMR, other professional service and humanoid segments	Section 10's frontier: 2 of 25 tasks clear the economic and capability gates at \$20 per human-equivalent hour, 9 at \$10, 16 at \$5. The base case assumes the industry works through the \$10-to-\$5 band by 2030 and no further.
Penetration beyond manufacturing	"Other professional service" units 98k → 950k; AMR 118k → 1.1m by 2040	Professional cleaning grew 34% and logistics service robots 14% in 2024, while hospitality fell 11% — penetration follows duty cycle, not sector glamour. ^[3]
Humanoid commercialisation	Units 12k (2025) → 600k (2030) → 6.0m (2040) in the base case, with ASP falling to \$13k	~19,100 units shipped globally in H1 2026, 97% Chinese. ^[33] The base path implies a 94.7% unit CAGR to 2030 from a tiny base — aggressive in growth terms, modest in absolute terms.
Labour scarcity and rising wages	Supports the unit paths in high-wage, ageing economies rather than entering as its own line	Section 13's caution applies: Japan's installations fell for a fourth consecutive year in 2024 while its working-age population shrank. ^{[1][21]} Demography raises the incentive; it does not deliver the order.
Economic competitiveness with human labour	The gate on every unit path	Section 9: a general-purpose platform at \$24.54 per human-equivalent hour clears \$29.11 US labour nominally but fails a 30% switching margin. Unit paths accelerate only where the margin is cleared.
Software, services and RaaS	Two base-driven segments: service at \$2,800→\$3,350 per unit-year; software at \$900→\$3,700	IFR reports service-robot RaaS growing 31% overall and 42% in logistics. ^[3] These lines convert a one-off sale into an annuity and are why the market's mix inverts by 2040.

M.3 Base case, 2026–2030

Exhibit 23 • Global robotics market by segment, 2025A–2030E, base case

Research question: how large is the robotics market today, how large could it be in 2030, and which layers grow? • **Units:** US\$ billion of annual industry revenue • **Period:** 2025 actual baseline, 2026–2030 modelled



Source: AUTHOR CALCULATION from the bottom-up baseline in M.1 and the drivers in M.2. Observed anchors from IFR unit and market-value data^{[1][3][52]}, A3 order statistics^{[42][43]}, Intuitive Surgical disclosure^[47] and third-party humanoid shipment estimates.^[33] **Illustrative scenario — not a market forecast.**

Takeaway: the base case takes the market from **\$107bn to \$193bn** between 2025 and 2030, a **12.4 per cent** compound rate — respectable, but not the exponential curve the category's narrative implies. More interesting is where the growth sits: of the \$85.5bn of incremental revenue, **software contributes 21 per cent, service and maintenance 19 per cent and humanoids 19 per cent**, while industrial robots and their integration together contribute 14 per cent. The industry grows by servicing and programming robots, not by selling more of them.

US\$ billion	2025A	2026E	2027E	2028E	2029E	2030E	CAGR 25–30	Share of growth
Industrial robots	17.5	18.5	19.7	20.9	22.1	23.5	6.1%	7.0%
Integration, peripherals and cells	35.0	36.1	37.3	38.5	39.8	41.1	3.3%	7.2%
AMR and logistics robots	5.3	6.1	6.9	7.9	9.1	10.4	14.3%	5.9%
Humanoid robots	0.6	1.2	2.3	4.4	8.6	16.8	94.7%	18.9%
Medical and surgical robots	14.5	16.2	18.0	19.9	22.0	24.2	10.8%	11.3%
Other professional service robots	4.3	4.9	5.6	6.4	7.3	8.3	13.9%	4.6%
Consumer service robots	8.2	8.9	9.6	10.3	11.1	11.9	7.6%	4.3%
Service, spares and maintenance	16.5	18.6	21.1	24.1	27.9	33.1	14.9%	19.4%
Robotics software, AI and simulation	5.3	7.0	9.3	12.5	16.9	23.5	34.7%	21.3%
Total	107.2	117.5	129.7	144.9	164.8	192.8	12.4%	100%
Year-on-year growth	—	9.6%	10.4%	11.7%	13.7%	17.0%		
Professional installed base, m units	5.9	6.5	7.2	8.1	9.2	10.7		

Exhibit 23b. **AUTHOR CALCULATION**. Note the acceleration in the growth rate, from 9.6% in 2026 to 17.0% in 2030: it is produced almost entirely by humanoids and software compounding off small bases, not by the mature segments improving.


M.4 Three scenarios, and what separates them

The scenarios differ in the variables the report has identified as decisive — task success, integration capacity, humanoid unit volume and software attach — not in a growth rate applied uniformly. Each is stated so that a reader can see which observation would move the industry from one path to another.

Scenario	What has to be true	2030	2035	2040	CAGR 25–30	CAGR 30–40	2040 vs 2025
Bear	Physical AI reliability plateaus below the level that clears a 30% switching margin outside structured environments. Integration capacity stays binding. Humanoids remain in extended pilots at 150k units in 2030. Software attach stalls at \$1,400 per unit-year.	\$144bn	\$200bn	\$278bn	6.1%	6.8%	2.6×
Base	The IFR trajectory holds through 2028 then moderates. Task success opens the tasks Section 10 places between \$10 and \$5 per human-equivalent hour. Humanoids commercialise in narrow logistics and manufacturing loops at 600k units in 2030. Software attach reaches \$2,200 per unit-year.	\$193bn	\$348bn	\$636bn	12.4%	12.7%	5.9×
Bull	Task success crosses the reliability threshold across several task families at once, opening the capability gate that closes 5 of 25 tasks in Section 10. RaaS removes the capex barrier for mid-sized buyers. Humanoids reach 1.5m units in 2030 on Chinese-scale supply chains. Software attach reaches \$2,900 per unit-year.	\$255bn	\$628bn	\$1.39tn	18.9%	18.4%	12.9×

Exhibit 24 • Scenario outcomes and the changing revenue mix

Research question: how wide is the plausible range for the robotics market, and how does its composition change? • **Units:** US\$ billion of annual industry revenue (left panel); share of revenue, per cent (right panel) • **Period:** 2025 actual; 2030, 2035 and 2040 modelled



Source: AUTHOR CALCULATION, drivers in M.2 and M.4. **Illustrative scenarios — not forecasts.** Long-horizon figures in particular are conditional arithmetic and should be read as a range, not a projection.

Takeaway (left): the 2030 range is narrow — \$144bn to \$255bn, a factor of 1.8 — because most of the 2030 market is already installed and under contract. By 2040 the range widens to a factor of five, from \$278bn to \$1.39 trillion, and essentially all of that spread is the humanoid and software assumption. **Takeaway (right):** the market's composition inverts. Integration falls from 33 to 6 per cent of revenue as standardisation does its work, while service, spares and software rise from 20 to 57 per cent. By 2040 the robotics industry earns most of its revenue from keeping robots working, not from selling them.

M.5 Which segments capture the growth

US\$ billion, base case	2025A	2030E	2035E	2040E	CAGR 25–40	Why it grows fast, or does not
Robotics software, AI and simulation	5.3	23.5	73.9	191.1	27.0%	The fastest and largest new pool. Charges per robot per year against the three variables Model 5 shows dominate cost: uptime, task success and relative productivity. Grows with the installed base, not with new sales.
Service, spares and maintenance	16.5	33.1	77.5	173.0	17.0%	Pure annuity on a base that grows nine-fold to 2040. The Investment Opportunity Map scored this theme highest for exactly this reason.
Humanoid robots	0.6	16.8	42.5	78.0	38.3%	Highest growth rate, smallest starting base, widest error bars. Note the shape: units rise 500× while ASP falls 74%, so revenue rises only 130×. Falling price is the mechanism of adoption and the brake on market value.
Other professional service robots	4.3	8.3	15.6	24.7	12.3%	Cleaning, agriculture, inspection and hospitality. Growth follows duty cycle: large-area cleaning compounds, hospitality does not.
AMR and logistics robots	5.3	10.4	16.9	24.2	10.6%	Strong unit growth (118k to 1.1m) largely offset by ASP falling from \$45k to \$22k. The most commoditised robot category.
Medical and surgical robots	14.5	24.2	36.8	51.8	8.9%	The most profitable segment and among the slowest-growing, because it is priced against clinical reimbursement rather than against labour cost.
Consumer service robots	8.2	11.9	17.1	22.6	7.0%	Enormous unit volume, trivial unit price. Section 3's point stands: the largest robot population in the world is economically the smallest market.
Industrial robots	17.5	23.5	27.8	31.2	3.9%	Mature. Units grow 2.3× to 2040 while ASP falls 21%. The category that defines the industry contributes least to its growth.
Integration, peripherals and cells	35.0	41.1	40.4	39.0	0.7%	The one segment that shrinks in real terms. The multiplier falls from 2.0× to 1.25× as standardisation does what Section 6 predicts — which is good for adoption and bad for integrators.
Total	107.2	192.8	348.4	635.6	12.6%	

Exhibit 25. Research question: which segments capture the growth, and why? Units: US\$ billion of annual industry revenue. Period: 2025 actual, 2030–2040 modelled.

AUTHOR CALCULATION. Segment shares of the 2030 market: integration 21.3%, service 17.2%, medical 12.5%, industrial robots 12.2%, software 12.2%, humanoids 8.7%, consumer 6.2%, AMR 5.4%, other professional 4.3%.

THE CROSS-CHECK THAT DISCIPLINES ALL OF THIS: REVENUE PER INSTALLED ROBOT

Divide the market by the professional installed base that generates it. In 2025 the industry earns **\$18,169 per installed professional robot per year**. By 2030 the base case has it at **\$18,034** — essentially unchanged — and by 2040 at **\$12,307** AUTHOR CALCULATION. The bear and bull cases converge on almost the same figure: \$12,498 and \$12,571 respectively.

That convergence is the most useful single result in this section. **The robotics market grows because there are more robots, not because each robot is worth more.** Revenue per unit falls in every scenario, because falling hardware prices and falling integration multiples outweigh rising software and service attach. Any market projection that assumes both rapid unit growth *and* stable revenue per unit is double-counting, and most published forecasts do exactly that.

M.6 How much of this is Physical AI?

The question can be answered directly by holding capability constant. The difference between the bear and base cases is almost entirely a Physical AI assumption: whether task success rises enough to open the tasks Section 10 prices between \$10 and \$5 per human-equivalent hour, and whether software can therefore charge \$2,200 per robot-year instead of \$1,400.

Attribution, base minus bear	2030	2040	Share of the gap
Total market gap	\$48.3bn	\$357.4bn	100%
Of which humanoid hardware	\$11.7bn	\$51.4bn	14%
Of which software, AI and simulation	\$10.8bn	\$144.3bn	40%
Of which service on a larger fleet	\$6.4bn	\$104.0bn	29%
Of which other hardware and integration	\$19.5bn	\$57.6bn	16%

AUTHOR CALCULATION. Roughly 70 per cent of the 2040 gap between the bear and base cases is software and the service revenue that a larger fleet generates — not robot hardware. Physical AI changes this market less by creating new machines than by making existing machines chargeable for more, and by enlarging the fleet those charges are levied on. That is the same conclusion Model 5 reached at the level of a single robot, arrived at independently at the level of an industry.

M.7 What would move the industry between scenarios

Five observations, in descending order of how much they move the model. Each is currently undisclosed, which is why the range is as wide as it is.

- **A published intervention rate in paid production.** It sets task success, which sets the frontier width in Section 10, which sets every unit path here. No vendor publishes one.
- **Humanoid units actually deployed rather than shipped.** One third-party estimate puts Unitree's active deployments near 250 against 5,500+ units shipped in 2025. ^[33] If that ratio persists, the base case's 600k units in 2030 is a shipment number and not a revenue-generating fleet.
- **Software attach rate.** Whether fleet software becomes a per-robot subscription or is bundled free into hardware decides a third of the long-run gap between scenarios.
- **The integration multiplier.** It falls from 2.0× to 1.25× in the base case. If standardisation does not arrive, the market is larger in dollars and smaller in robots — a worse outcome for everyone except integrators.
- **Actuator, magnet and grid capacity.** Section 20's constraint migration binds the bull case physically: 15m humanoids a year in 2040 requires magnet capacity that does not exist outside China and is not currently being built. ^{[22][23]}

THE FIVE ANSWERS THIS SECTION SET OUT TO GIVE

How large is the robotics market today? About **\$107bn** in 2025, built bottom-up. The robots themselves are \$50bn of that; the rest is integration, service and software.

How large could it be in 2030? **\$144bn to \$255bn**, with a base case of **\$193bn** — a 12.4 per cent compound rate. The near-term range is narrow because most of the 2030 market is already installed.

What about 2035-2040? **\$278bn to \$1.39 trillion** by 2040, base case **\$636bn**. The range widens five-fold because the long-run answer depends almost entirely on whether humanoids and fleet software become real businesses.

What creates the growth? Not robot sales. Software (21 per cent of incremental revenue to 2030), service and maintenance (19 per cent) and humanoids (19 per cent) — against 14 per cent from industrial robots and their integration combined.

How much does Physical AI change the curve? It is the difference between a 6 per cent industry and a 12-19 per cent one, and roughly seventy per cent of that difference arrives as software and service revenue rather than as machines. Physical AI does not principally make the robotics market bigger by adding robots; it makes it bigger by making each robot worth paying for, every year, for longer.

INVESTMENT OPPORTUNITY MAP**Capital-allocation themes across the robotics value chain****SCOPE AND LIMITS OF THIS SECTION**

This section identifies **capital-allocation themes** and the diligence questions that separate a real business from a demonstration. It contains **no security recommendations, no company rankings, no price targets, no trading advice and no expected returns**. Companies appear only as evidence for or against a theme, with the status of each claim stated. The preceding section sizes the industry these themes sit inside; it values a market, not a security. Nothing here is investment advice; this report is not a substitute for professional advice, and readers should form their own view and consult a qualified adviser before acting.

The framework

AUTHOR FRAMEWORK — OPPORTUNITY QUALITY

$$\text{Opportunity quality} = \frac{\text{bottleneck intensity} \times \text{deployment evidence} \times \text{recurring-revenue potential}}{\text{capital intensity} \times \text{technical risk} \times \text{adoption friction}}$$

AUTHOR FRAMEWORK . Each factor is scored 1-5 in the direction of its label — higher numerator is better, higher denominator is worse. The output is an ordinal comparison across themes, not a valuation or an expected return. Its purpose is to make the trade-off between "solves a real bottleneck" and "is expensive and risky to build" explicit.

Theme	Bottle-neck	De-ploy. evid.	Recur. rev.	Capit-al in-tens.	Tech. risk	Ad-opt. fric-tion	Score	Horizon	Bottleneck solved · business model · principal failure signal
Integration & deployment tooling	5	5	3	2	2	2	9.4	Near	Bottleneck: integration is 12–60% of installed capital and the reason US density is a quarter of Korea's at similar wages. ^[2] Model: software licences, configuration tools, services attach. Recurring: moderate — some per-seat and per-cell licensing. Failure signal: commissioning time per cell not falling across a customer cohort; revenue mix drifting back toward bodyshop services.
Maintenance, spares & field service	4	5	5	2	1	2	25.0	Near	Bottleneck: maintenance is 15.6% of five-year TCO and downtime destroys the denominator. Model: service contracts, spares, uptime SLAs. Recurring: highest of any theme — it is annuity revenue on an installed base of 4.66m machines. Failure signal: attach rate falling as OEMs internalise service; margin compressing on parts as designs standardise.
Robots-as-a-Service & fleet finance	4	4	5	5	3	2	2.7	Near	Bottleneck: converts an irreversible capex decision into an opex one, which is exactly the switching-margin problem in Model 2. Evidence: IFR reports service-robot RaaS up 31%, and 42% in transport and logistics. ^[3] Model: per-hour or monthly rate. Capital intensity is the risk: the provider owns the asset and the residual. Failure signal: utilization below underwriting, residual-value writedowns, contract non-renewal, or funding cost rising faster than rate cards.
Actuators, reducers, drives, encoders	5	5	2	4	2	2	3.1	Near-medium	Bottleneck: the largest BOM block and the least deflationary; precision reducers remain concentrated. Model: component sales, design-in, long qualification cycles. Recurring: low — unit sales, though spares help. Failure signal: customer in-sourcing of joint modules; Chinese entrants taking design-in slots at qualified quality.
Simulation, data & fleet software	4	4	5	2	3	3	4.4	Medium	Bottleneck: the U , Q and ϕ terms — the three largest operational levers in Model 5. Evidence: DeepFleet claimed at 10% fleet travel-time improvement ^[27] ; $\pi_{0.5}$ scaling in environments ^[40] . Model: subscription; the buyer's S term. Failure signal: customers unable to attribute measured uptime or task-success improvement to the software; foundation-model providers commoditising the layer from above.
Safety, testing & regulatory compliance	4	3	4	2	2	3	4.0	Medium	Bottleneck: every model update to a machine near people is a safety-relevant change; five of 25 screened tasks are blocked on non-cost grounds. Model: certification, test tooling, continuous-validation software. Recurring: good — validation repeats with every release. Failure signal: standards bodies converging on a self-certification regime that removes the third-party need.
Vertical task robotics	3	4	4	4	3	3	1.3	Near-medium	Bottleneck: a specific high- λ task with a proven buyer — the tasks on the right of Exhibit 10. Evidence: professional cleaning +34%, logistics service robots +14% in 2024. ^[3] Model: system sale or RaaS. Failure signal: a general-purpose platform reaching adequate ϕ on the same task, or the vertical proving too small to amortise its own software.
Energy, charging & facility infrastructure	4	3	4	4	2	3	2.0	Medium-long	Bottleneck: charging caps the operating window (16h per charge for Digit v4 ^[29]) and interconnection caps fleet scale; grid investment must rise ~50%. ^[23] Model: equipment plus service contracts. Failure signal: battery energy density improving fast enough to make charging infrastructure a commodity.
Permanent magnets & critical-material diversification	5	3	2	5	3	4	0.5	Long	Bottleneck: the most concentrated dependency in the chain — China at 94% of sintered magnets, ex-China capacity under 20% of 2035 demand, ~\$60bn of investment required. ^[22] Model: long-term offtake, price floors, state support. Failure signal: the one to watch closely — China's November 2025 suspension of controls shows how fast the policy-created price umbrella can be withdrawn, which is precisely how earlier Western rare-earth projects failed.

Exhibit 20. Research question: where in the value chain does solving a real bottleneck coincide with a durable business model? Units: ordinal 1–5 scores; score = product of numerator factors divided by product of denominator factors. Period: assessed August 2026. Evidence status: **AUTHOR FRAMEWORK** throughout; cited figures as tagged. **Takeaway:** the highest-scoring themes are the least glamorous. Maintenance and field service scores highest because it has annuity economics against an installed base that already exists; permanent magnets scores lowest not because the bottleneck is small — it is the largest in the report — but because it is capital-intensive, slow and policy-exposed. A high bottleneck score is not a business.

Capital sequencing by evidence horizon

Horizon	What is already proven	What must be proven next	Themes that fit
Near evidence exists today	4.66m operating industrial robots; 542,076 installations a year; RaaS growing 31%; a million-robot single-operator fleet ^{[1][3][27]}	Nothing — these are operating businesses. The question is competitive position, not technical feasibility.	Maintenance and field service; integration tooling; RaaS in logistics; vertical task robotics in cleaning and transport
Medium pilots exist, unit economics un-proven	Paid humanoid deployments at 9 sites; 65,000+ cumulative operating hours; 100,000+ totes ^{[29][30]}	An intervention rate. A published measure of human interventions per task hour in paid production, and its trend. No public source provides one.	Fleet software and data; safety and validation; general-purpose platforms in high-wage two-shift settings
Long structural, slow, policy-exposed	The scale of the constraint: 94% magnet concentration; grid investment gap; integrator shortage ^{[22][23]}	That demand at scale actually arrives — and that policy support survives a change of government and a Chinese price response.	Magnets and critical materials; actuator and reducer capacity; charging and grid infrastructure

Exhibit 21. Research question: in what order does evidence become available, and what should be underwritten at each stage? Units: qualitative. Period: assessed August 2026.

AUTHOR FRAMEWORK

Project-level diligence gates

Eleven questions. A project that cannot answer the first six with numbers from paid production is a demonstration, whatever its capability.

#	Gate	What good looks like, and why it matters
1	Paid commercial deployment	Revenue from a customer who is not an investor, a strategic partner or a pilot grant. Rung 4 or above on the Section 4.1 ladder.
2	Productive-hour data	Scheduled, operating and productive hours reported separately. BMW's 1,250 operating hours is the public benchmark; almost no vendor matches that disclosure. ^[25]
3	Intervention rate	Human interventions per task hour, and its trend over the deployment. The single most diagnostic number in robotics, and the one least often disclosed.
4	Repeat orders and renewals	Second and third orders from the same customer, and RaaS renewal rates. A first order proves curiosity; a third proves economics.
5	Integration time	Days from delivery to production, and whether it is falling across a customer cohort. This is the <i>I</i> term and the scaling constraint.
6	Gross margin after service cost	Margin with field service, spares, teleoperation and customer engineering fully loaded — not hardware gross margin. The teleoperation trap of Section 17.2 shows up here first.
7	Warranty reserves	Reserve as a percentage of revenue, and its trend. A rising reserve on a maturing product is a reliability disclosure by another name.
8	Customer concentration	Revenue share of the top customer. A single anchor customer who is also an investor or supplier is not market validation.
9	Safety performance	Recordable incidents per operating hour, certification status, and how model updates are validated before release.
10	Realistic utilization	Underwritten <i>H</i> , <i>U</i> and <i>Q</i> compared with what the deployment actually achieves. An underwriting assumption of 0.90 uptime against an industry-observed 0.58 is the most common error in the sector.
11	Path to positive unit economics	Which of Model 5's eleven variables must move, by how much, by when — and which of them the company controls. A plan that depends on all four levers moving at once is a plan with four independent failure modes.

Company evidence register

Exhibit 22. Companies referenced in this report, with the status of each claim. Research question: what does each company's disclosure actually establish? Units: qualitative. Period: evidence to August 2026. Used as evidence only, never as a recommendation.

Company	Claim used in this report	Status	What it does not establish
Amazon	One millionth robot deployed July 2025, 300+ facilities; DeepFleet claimed 10% fleet travel-time improvement	COMPANY-REPORTED	That the model works in buildings Amazon does not design. No cost-per-productive-hour disclosed.
Amazon	Blue Jay handled ~75% of item types; withdrawn from operations February 2026	COMPANY-REPORTED	Why it was withdrawn. The disclosure is unusually valuable precisely because negative results are rarely published.
Amazon	Reported internal plans to avoid ~600,000 future US hires; \$2–4bn/yr savings from 2027 (Morgan Stanley), up to \$10bn internally	THIRD-PARTY ESTIMATE of leaked internal planning	Anything about actual deployments. These are plans reported by journalists, not results.
BMW & Figure	Figure 02, ~1,250 operating hours over ten months, 90,000+ components, 1.2m steps, 30,000+ X3 vehicles; described as a pilot. Figure 03 to follow in logistics sequencing	COMPANY-REPORTED , published by the customer	Cost, intervention rate or profitability. BMW calls it a pilot project, and the report treats it as one.
Agility Robotics & GXO	100,000+ totes; 65,000+ cumulative operating hours; 9 sites; \$125,000 BOM for Digit v4; 16h per charge; RoboFab up to 10,000 units/yr; \$300m+ pre-orders; \$620m SPAC at \$2.5bn pre-money	COMPANY-REPORTED ; capacity and pre-orders are ANNOUNCED	Unit economics. Tote handling is one bounded loop. Pre-orders are commitments, not revenue.
Unitree	G1 from \$13,500; 2025 revenue ¥1.708bn, 62.9% humanoid gross margin; ~5,900 units shipped H1 2026; ~250 estimated active deployments from 5,500+ shipped in 2025	COMPANY-REPORTED financials; shipments and deployments THIRD-PARTY ESTIMATE	That the machines are performing paid labor. A list price is not a labor cost, and shipped is not deployed.
Tesla	Optimus V3 reveal moved to mid-2026; Fremont line conversion; no 2026 volume commitment; ~10,000 unique parts	ANNOUNCED TARGET	Any production. The report's clearest case of the ambition/commercialisation distinction.
Intuitive Surgical	11,106 da Vinci systems at end-2025, +12%; Q4 procedures +17%; FY revenue \$10.06bn	COMPANY-REPORTED , audited financials	Labor substitution. This is a clinician-augmentation business, and the closest thing robotics has to a proven long-run service model.
NVIDIA	Jetson AGX Thor: 2,070 FP4 TFLOPS, 128GB, 130W, \$3,499 dev kit; 7.5× Orin compute; named adopters	COMPANY-REPORTED	That edge compute is the constraint. Model 5 ranks compute-related cost well below utilization and productivity.
FANUC, ABB, KUKA, Yaskawa	Represented in the IFR installation and stock series and the A3 order data rather than by company disclosure	OBSERVED DATA at industry level	Individual company positions. The report deliberately uses industry aggregates for the mature segment.

APPENDICES

Appendix A • Model definitions and assumption register

Every input used in Models 1–5, with its value, its evidence status and the reason it was chosen. Inputs marked **RESEARCH ASSUMPTION** are choices made to expose sensitivity, not estimates of the world; they are all swept in Model 5.

A.1 Formulae

THE COMPLETE MODEL SET

$$\mathbf{M1} \quad \text{CRF}(r,n) = r(1+r)^n / [(1+r)^n - 1]$$

$$C_R = [\text{CRF}(r,n)(P+I) + M + S + O] / (H U Q) + E + V ; C_{R,HE} = C_R / \varphi$$

$$\mathbf{M2} \quad P^* = [(\varphi W - E - V)(H U Q) - (S+O)] / [\text{CRF}(r,n)(1+i) + m]$$

$$\mathbf{M3} \quad C_{\text{task}} = L(1/\lambda)x ; \text{enters} \Leftrightarrow C_{\text{task}} < W_{FL}/1.30 \text{ and gate} \leq 2 ; L^* = (W_{FL}/1.30)(\lambda/x)$$

$$\mathbf{M4} \quad \text{M1 evaluated across } H \in \{8,16,20,24\} \text{ h/day with regime-specific } U, Q$$

$$\mathbf{M5} \quad \text{M1 evaluated over each input's tested range, one at a time and in combination}$$

A.2 Input register

Input	Case A	Case B	Case C	Status	Basis and Model 5 sweep range
Purchase price P	\$75,000	\$75,000	\$75,000	ASSUMPTION	Held identical across cases so the three differ only in use. Sweep \$20,000–\$150,000. Anchors: Agility's \$125,000 disclosed BOM ^[29] , Unitree G1 from \$13,500 ^[32] ; A3 average order value \$61,197 per unit in 2025 ^[43] .
Integration I , % of P	150%	35%	35%	ASSUMPTION	Industrial cells historically 1–2× hardware ^[44] ; the general-purpose value proposition is a much lower figure. Sweep 10%–150%.
Financing r	8%	10%	10%	ASSUMPTION	Mature vs emerging asset class. Sweep 3%–18%.
Useful life n	8 yr	5 yr	4 yr	ASSUMPTION	Obsolescence-bounded on fast-moving platforms. Sweep 3–12 years.
Maintenance M , % of P /yr	6%	12%	15%	ASSUMPTION	Scales with actuated joint count, duty cycle and design maturity. Sweep 3%–25%.
Software S	\$3,000	\$9,000	\$12,000	ASSUMPTION	Fleet management, inference and controls subscription. No vendor publishes list pricing. Sweep \$1,000–\$24,000/yr.
Insurance/other O	\$3,000	\$2,000	\$2,000	ASSUMPTION	Guarded industrial cells carry higher fixed compliance cost; mobile platforms carry higher liability exposure. Not swept.
Scheduled H	5,280 h	5,280 h	2,500 h	ASSUMPTION / CALIBRATED	16h × 330d for A and B; Case C mirrors BMW's stated 10-hour weekday pattern. ^[25] Sweep 8–24 h/day.
Uptime U	0.92	0.85	0.58	ASSUMPTION / CALIBRATED	Case C from ~1,250 operating hours against ~2,150 scheduled. ^[25] Sweep 0.50–0.97.
Task success Q	0.95	0.75	0.75	ASSUMPTION	The report's weakest input. No public source discloses task-success or intervention rates in paid production. Sweep 0.50–0.97.
Energy E	\$0.60/h	\$0.30/h	\$0.30/h	CALCULATION	~3 kW and ~1.5 kW average draw at \$0.20/kWh industrial tariff. Not swept — immaterial at the unit level.
Supervision V	\$1.80/h	\$3.00/h	\$6.00/h	CALCULATION	\$45/h fully loaded technician (consistent with the \$27.89 mean wage for industrial machinery mechanics ^[10] at a 1.55× loading, rounded up for shift premium) at 1:25, 1:15 and 1:7.5 ratios. Sweep \$0.25–\$8.00/h.
Relative productivity φ	1.50	0.70	0.70	ASSUMPTION	Purpose-built cells exceed human throughput on their single task; general platforms do not yet match it. Sweep 0.40–1.40.
Operating days	330	330	250	ASSUMPTION	Case C uses weekdays only, per the BMW pattern.
Switching margin k	1.30 throughout			ASSUMPTION	Firms require a cost advantage, not equality, before an irreversible capital commitment. Broadly equivalent to a two-year payback requirement on the incremental investment.

A.3 Exclusions and known limitations

- **Residual value is set to zero.** Any positive resale value would lower every cost figure in the report. This makes all robot costs conservative (high).
- **Taxes, depreciation allowances and investment incentives are excluded.** These vary by jurisdiction and would generally favour the robot.
- **Facility cost beyond the integration percentage is excluded** — floor space, utilities capacity, network. This makes robot costs optimistic (low).
- **No learning curve is applied within a deployment.** In practice U and Q improve with tenure; the model holds them constant, which is conservative for a mature deployment and generous for a new one.
- **The Q and φ terms are unobservable from public data.** This is the report's principal limitation and the reason three cases are presented rather than one.
- **The human benchmark is US-specific.** Applying it to other labour markets requires substituting local compensation data; the loading factor of 1.551× is derived from US BLS ratios and will differ elsewhere, materially so where statutory social contributions are higher.

Appendix B • Automation-frontier task screen

The complete 25-task screen behind Model 3 and Exhibits 10, 10b and 10c. λ is duty-cycle attainment, $\times$ the exception-burden multiplier, and L^* the break-even reference platform cost after a 30 per cent switching margin.

Sector	Task	Median wage	Loading	Fully loaded	Threshold (+1.30)	λ	$\times$	Capability gate	L^*	Enters at
Manufacturing	Spot welding, body-in-white	\$24.95	1.44×	\$35.93	\$27.64	1.00	1.05	1 — passed	\$26.32	\$20
Manufacturing	Structured visual inspection	\$19.85	1.44×	\$28.58	\$21.99	1.00	1.05	1 — passed	\$20.94	\$20
Manufacturing	Machine tending, CNC load/unload	\$19.85	1.44×	\$28.58	\$21.99	0.95	1.10	1 — passed	\$18.99	\$10
Manufacturing	Palletising and depalletising	\$18.77	1.44×	\$27.03	\$20.79	0.90	1.10	1 — passed	\$17.01	\$10
Warehousing	Tote and container movement	\$18.77	1.43×	\$26.84	\$20.65	0.85	1.15	1 — passed	\$15.26	\$10
Healthcare	Supply and linen transport	\$20.32	1.42×	\$28.85	\$22.20	0.70	1.20	1 — passed	\$12.95	\$10
Manufacturing	Mixed-model dexterous final assembly	\$19.85	1.44×	\$28.58	\$21.99	0.90	1.60	3 — uncertain	\$12.37	never
Warehousing	Packing and carton forming	\$17.53	1.43×	\$25.07	\$19.28	0.80	1.25	2 — near	\$12.34	\$10
Warehousing	Each-picking from mixed bins	\$19.05	1.43×	\$27.24	\$20.95	0.80	1.50	2 — near	\$11.17	\$10
Retail	Shelf scanning and inventory audit	\$19.05	1.35×	\$25.72	\$19.78	0.60	1.10	1 — passed	\$10.79	\$10
Cleaning	Floor scrubbing, large open areas	\$17.71	1.35×	\$23.91	\$18.39	0.50	1.15	1 — passed	\$8.00	\$5
Agriculture	Milking	\$15.81	1.32×	\$20.87	\$16.05	0.55	1.15	1 — passed	\$7.68	\$5
Construction	On-site material transport	\$18.89	1.48×	\$27.96	\$21.51	0.50	1.40	2 — near	\$7.68	\$5
Retail	Back-room replenishment	\$19.05	1.35×	\$25.72	\$19.78	0.50	1.30	2 — near	\$7.61	\$5
Warehousing	Floor-loaded trailer unloading	\$18.77	1.43×	\$26.84	\$20.65	0.55	1.50	2 — near	\$7.57	\$5
Construction	Rebar tying and layout marking	\$18.89	1.48×	\$27.96	\$21.51	0.40	1.50	2 — near	\$5.74	\$5
Construction	Finish trades in occupied space	\$24.00	1.48×	\$35.52	\$27.32	0.40	2.00	4 — blocked	\$5.46	never
Hospitality	Short-order fry and beverage station	\$16.98	1.30×	\$22.07	\$16.98	0.40	1.30	2 — near	\$5.22	\$5
Hospitality	Food running and bussing	\$16.98	1.30×	\$22.07	\$16.98	0.35	1.40	2 — near	\$4.25	\$2
Healthcare	Patient lifting and transfer	\$20.32	1.42×	\$28.85	\$22.20	0.30	1.60	4 — blocked	\$4.16	never
Cleaning	Restroom and detail cleaning	\$17.71	1.35×	\$23.91	\$18.39	0.40	1.80	3 — uncertain	\$4.09	never
Agriculture	Field spraying and mechanical weeding	\$15.81	1.32×	\$20.87	\$16.05	0.25	1.30	1 — passed	\$3.09	\$2
Agriculture	Delicate fruit harvesting	\$15.81	1.32×	\$20.87	\$16.05	0.15	1.80	3 — uncertain	\$1.34	never
Home services	Laundry folding, home setting	\$17.21	1.00×	\$17.21	\$13.24	0.10	1.50	3 — uncertain	\$0.88	never
Home services	General household tidying	\$17.21	1.00×	\$17.21	\$13.24	0.10	2.00	4 — blocked	\$0.66	never

Wages: BLS OEWS May 2025 national medians.^[10] Loading factors: 1.551× for standard employed occupations, adjusted downward for sectors with lower statutory burden and set to 1.00 for household work, where the counterfactual is unpaid or informally paid labour. λ , $\times$ and gate assignments are **AUTHOR FRAMEWORK**; L^* is an **AUTHOR CALCULATION**. "Enters at" shows the highest tested reference platform cost at which both gates pass — that is, the level at which the task first becomes automatable. "never" means the task does not enter at any tested level, including \$1 per hour.

Appendix C • Human and robot cost normalisation

C.1 Fully loaded human cost by wage level

Wage	Benefits & payroll	ECEC total compensation	Recruiting & turnover	Absence backfill	Injury loading	Fully loaded employer cost
\$15.81	\$6.81	\$22.62	\$0.95	\$0.63	\$0.32	\$24.52
\$17.71	\$7.63	\$25.34	\$1.06	\$0.71	\$0.35	\$27.46
\$18.77	\$8.08	\$26.85	\$1.13	\$0.75	\$0.38	\$29.11
\$19.85	\$8.55	\$28.40	\$1.19	\$0.79	\$0.40	\$30.78
\$24.95	\$10.74	\$35.69	\$1.50	\$1.00	\$0.50	\$38.69
\$32.60	\$14.04	\$46.64	\$1.96	\$1.30	\$0.65	\$50.55

Wage column: BLS OEWS May 2025 medians for crop farmworkers, janitors, hand laborers/material movers, team assemblers, welders, and the US private-industry average wage respectively.^[10] Benefits derived from the BLS ECEC March 2026 wage share of 69.9 per cent of private-industry compensation.^[9] The last three columns are

RESEARCH ASSUMPTIONS at 6%, 4% and 2% of wage. **AUTHOR CALCULATION**. Constant loading factor 1.551x.

C.2 Robot cost per human-equivalent hour, by purchase price and configuration

Purchase price	Case A — specialised industrial cell	Case B — general-purpose platform (mature)	Case C — pilot-observed calibration
\$150,000	\$13.19	\$39.69	\$140.87
\$100,000	\$9.62	\$29.59	\$103.04
\$75,000	\$7.83	\$24.54	\$84.13
\$50,000	\$6.04	\$19.49	\$65.22
\$30,000	\$4.61	\$15.45	\$50.09
\$20,000	\$3.90	\$13.42	\$42.52

AUTHOR CALCULATION, Model 1 with all inputs other than P held at the values in Appendix A.2. **Takeaway:** the three configurations never overlap at any price. A \$150,000 machine used like Case A (\$13.19/h) beats a \$20,000 machine used like Case C (\$42.52/h) by more than three to one. How the machine is used dominates what it costs.

Appendix D • Contradiction register

Eleven places where sources disagree, or where a widely repeated figure does not survive checking. Each is shown with both positions and the reason for the disagreement, rather than resolved by preference.

#	Contradiction	The two positions	Resolution used in this report
C1	China's robot density	IFR WR2024: 470 per 10,000 manufacturing employees for 2023, ranking China third in the world. IFR WR2025: 166 for 2024, ranking China outside the top ten. ^{[2][4]}	Not a data error and not a fall in robots: China's National Bureau of Statistics revised the manufacturing employment denominator. Both vintages are used, labelled by reference year, and never compared as a time series. The global average moved from 162 to 132 for the same reason.
C2	China robot production vs world installations	China Machinery Industry Federation: 595,000 industrial robots produced in 9M 2025. ^[45] IFR: 542,076 installed worldwide in all of 2024. ^[1]	Different definitions and different concepts. The Chinese statistic counts manufactured units on a national definition that includes simpler arms and manipulators; the IFR counts installations meeting ISO 8373. Production ≠ installation; national definition ≠ IFR definition. Never combined.
C3	Robot prices: falling or rising?	IFR/ARK: robot prices fell ~80% 1995–2017 and 83% quality-adjusted 1990–2005. ^[44] A3: average value per North American robot ordered rose from \$61,197 (2025) to \$64,796 (H1 2026). ^{[42][43]}	Both true, measuring different things. The historical series is quality-adjusted over three decades; the A3 series is nominal average order value over eighteen months and reflects a mix shift toward higher-content systems. The report declines to state a current robot learning rate and argues the correct denominator is delivered productive work.
C4	Robots and employment	Acemoglu & Restrepo: –0.18 to –0.34pp employment-to-population per robot per thousand workers. ^[12] Aghion et al.: firm-level employment elasticity +0.2 to +0.4 . ^[15] Dauth et al.: manufacturing losses fully offset by services. ^[14]	Different units of observation and different equilibrium scope. Commuting-zone designs net out national offsets by construction; firm-level designs capture business-stealing from non-adopters; German results embed institutions that redeploy internally. All three are reported; none is used alone.
C5	Robots and productivity	Graetz & Michaels: robots added 0.36pp to annual labour productivity growth, 1993–2007. ^[13] NY Fed/BLS: US manufacturing labour productivity growth was –0.5%/yr , 2010–2022, while the global robot stock nearly tripled. ^[11]	Different periods, geographies and levels of aggregation. The report treats the post-2010 US series as an unexplained counterexample that any robot-productivity thesis must address, and does not claim to resolve it.
C6	Humanoid units: shipped, deployed or working?	~19,100 humanoid units shipped globally in H1 2026, 97% Chinese. ^[33] The same source estimates ~250 Unitree units in active deployment from 5,500+ shipped in 2025.	Shipped, deployed and operating productively are three different states. The report reports shipments as third-party estimates and treats deployment claims separately, requiring rung 4 evidence before treating any unit as performing labour.
C7	Humanoid utilization assumptions	Vendor and analyst cost models commonly assume 0.90+ uptime and near-continuous operation. BMW's published record implies ~0.58 operating hours to scheduled hours. ^[25]	The report presents both: Case B at 0.85 as a mature-operation assumption, Case C at 0.58 as the pilot-observed calibration. The \$24.54 vs \$84.13 gap is the size of this disagreement, and it is larger than any hardware price debate.
C8	Amazon: hiring or replacing?	Amazon: 250,000 seasonal roles for the 2025 holiday, 700,000+ employees upskilled since 2019, 30% more technical roles at its newest centre. ^{[27][28]} Reported internal planning: avoid ~600,000 future US hires; ~30,000 corporate role reductions. ^[48]	These are compatible: not-hiring in future is different from firing now, and both can be true simultaneously in different parts of one firm. The report labels the first company-reported and the second a third-party account of internal planning, and treats neither as an observed employment outcome.
C9	Blue Jay: capability success or commercial failure?	Amazon: Blue Jay handles ~75% of item types and consolidates three robotic stations into one. Amazon, February 2026: "no longer utilizing Blue Jay in operations." ^[28]	Both are company-reported and both stand. Technical capability and operational viability are different tests, and this is the clearest public example of a system passing the first and failing the second. Used in the report as evidence for the Section 17 link-5 gap.
C10	Demographics and robot demand	Working-age population is falling in Japan (~16.6% from peak), Germany and Italy. ^[21] Japanese robot installations fell for a fourth consecutive year in 2024; German installations fell 5%. ^[1]	No contradiction in theory, but a direct contradiction of the popular inference. Demography raises the long-run incentive; the capital cycle, end-market conditions and integrator availability dominate over a five-year horizon. The report treats link 5 of the demographic chain as the weak one.
C11	Service-robot counts across report editions	Year-on-year growth rates are widely quoted from IFR service-robot data. The IFR states the data is <i>not projected to the whole industry</i> , that the 294-supplier sample composition varies each year, and that comparison across report editions is discouraged. ^[3]	Service-robot figures are used only within a single report edition, are labelled "observed sample" rather than "observed data", and are never added to industrial installations.

Evidence gaps — things no source discloses

Missing number	Why its absence matters
Intervention rate in paid production	The most diagnostic operational metric in robotics. Not published by any vendor reviewed. Model 1's Q term is therefore an assumption in every case.
Cost per productive hour, disclosed by an operator	No public deployment discloses one. Every figure in Sections 7–12 is therefore modelled, not observed.
RaaS rate cards	Providers describe per-hour and monthly pricing but do not publish rates. This removes the most direct available market check on Model 1.
Quality-adjusted robot price series post-2005	Without it, no defensible learning rate can be computed for the modern industry.
Actuator and precision-reducer capacity data	The largest BOM block has no public capacity series, so the supply constraint cannot be sized.

Appendix E • Thesis scorecard

Eight sub-claims, each scored against the evidence assembled in this report. The overall verdict is derived from these scores rather than asserted independently.

#	Sub-claim	Verdict	Decisive evidence
1	Robots already deliver machine labour below human cost in some tasks	Supported	Case A at \$7.83 per human-equivalent hour against \$29–39 fully loaded human cost, with 4.66m operating units and 542,076 annual installations as the revealed preference. ^[1]
2	Physical AI is widening the set of economically automatable tasks	Supported	Open X-Embodiment cross-embodiment gains ^[38] ; $\pi_{0.5}$ generalising to unseen homes and scaling in environment count ^[40] ; Gemini Robotics 1.5 motion transfer across embodiments ^[36] . Benchmarks, but consistent and independent.
3	General-purpose platforms are at cost parity with human labour today	Rejected	\$24.54 against \$29.11 fails a 30% switching margin; the only multi-month operating record implies \$84.13. ^[25] BMW itself calls its deployment a pilot.
4	Robot hardware is on a battery- or solar-like cost curve	Rejected	Average value per North American robot ordered rose 5.9% into H1 2026. ^{[42][43]} No quality-adjusted price series exists for the last two decades. The historical 1990–2005 decline is real and thirty years old.
5	Utilization and productivity matter more than hardware price	Supported	Model 5: ϕ range \$30.67 and H range \$26.43 against \$26.27 for a full price sweep; three of the top five variables are operational rather than hardware.
6	Falling robot cost would open a broad automation frontier	Partially supported	Model 3: 2 tasks at \$20 rising to 18 at \$2 — a real widening. But the count saturates at 18 of 25, with five tasks cost-viable at \$1 and blocked on non-cost grounds.
7	Robotics is a general deflationary force in consumer prices	Rejected	Exhibit 17: maximum modelled net effect of 5.2% in logistics, 0.3% in hospital care. Aghion et al. find producer prices -0.05 against profits $+0.78$ — the gain is retained, not passed on. ^[15]
8	If machine labour scales, new constraints bind	Supported	China at 94% of sintered magnets with ex-China capacity under 20% of 2035 demand ^[22] ; electricity demand growing 3.6%/yr to 2030 with grid investment needing to rise ~50% ^[23] ; US robot density one-quarter of Korea's at similar wages, indicating a deployment-capacity constraint. ^[2]

SCORE SUMMARY AND OVERALL VERDICT

4 supported • 1 partially supported • 3 rejected. The supported claims concern what machine labour already does and what constrains it. The rejected claims are all versions of the strong thesis: that general physical labour is now cheap, that hardware is deflating like a commodity, and that consumers will see it in prices.

Overall: PARTIALLY SUPPORTED. Robotics is converting selected physical work from a scarce human input into a scalable capital service, and Physical AI is widening the selection. It is not making physical labour cheap in general, and the mechanism that would be required for it to do so — near-zero marginal cost of an additional unit of capacity — does not exist for machines the way it does for models.

Appendix F • Exhibit register

#	Exhibit	Type	Evidence status	Source
1	Global industrial robot installations, 2014–2024 + outlook	Chart	OBSERVED + ANNOUNCED	IFR World Robotics 2022–2025 ^{[1][5][6][7]}
2	Global operational robot stock, 2015–2024, with China share	Chart	OBSERVED	IFR ^{[1][5][6][7][46]}
3	Robot density by country and the 2024 methodology revision	Chart	OBSERVED	IFR ^{[2][4]}
4	Task-level automation readiness framework	Table	FRAMEWORK + CALCULATION	BLS ^{[9][10]} ; author scoring
5	Commercial deployment evidence ladder	Table	Mixed, tagged per row	BMW, Amazon, Agility, Tesla, Intuitive ^{[25][27][28][29][34][47]}
6	Robot total-cost stack, \$/productive hour, three cases	Chart	CALCULATION	Model 1
7	Comparative technology price declines, annualised	Chart	OBSERVED + 3RD-PARTY	BNEF ^[24] ; IFR via Construction Physics ^[44] ; A3 ^{[42][43]}
8	Model 1 output: cost per productive and human-equivalent hour	Chart	CALCULATION	Model 1; BLS benchmark ^{[9][10]}
9	Human–robot parity matrix	Chart	CALCULATION	Model 2
9b	Break-even purchase price by wage and operating regime	Table	CALCULATION	Model 2
10	Automation frontier: break-even reference cost by task	Chart	CALCULATION + FRAMEWORK	Model 3; BLS ^{[9][10]}
10b	Frontier width at \$20, \$10, \$5, \$2 and \$1	Chart	CALCULATION	Model 3
10c	Frontier by sector	Table	CALCULATION	Model 3
11	Utilization economics: scheduled, operating, productive hours	Chart	CALCULATION	Model 4
11b	Utilization regimes, full table	Table	CALCULATION	Model 4
12	Robot cost sensitivity matrix (tornado)	Chart	CALCULATION	Model 5
12b	Sensitivity ranking, full table	Table	CALCULATION	Model 5
12c	Lever combinations required to reach \$10 and \$5	Chart	CALCULATION	Model 5
13	Productivity–literature contradiction matrix	Table	OBSERVED / peer-reviewed	Six studies ^{[11][12][13][14][15][16]}
14	Working-age population, indexed to 2000	Chart	OBSERVED	World Bank WD ^[21]
15	China–United States robotics ecosystem comparison	Table	Mixed FRAMEWORK +	IFR, IEA, NVIDIA, Amazon, Agility ^{[1][2][22][27][29][41][45][46]}
16	The robotics flywheel and its failure points	Diagram	FRAMEWORK	Author, with per-link evidence
16b	Flywheel link-by-link audit	Table	FRAMEWORK	Author, with per-link evidence
17	Cost-to-consumer-price transmission by sector	Chart	ASSUMPTION + CALCULATION	Author scenario
18	Displacement / productivity / reinstatement framework	Table	OBSERVED / peer-reviewed	Acemoglu & Restrepo, Graetz & Michaels, Dauth et al., Amazon ^{[12][13][14][27]}
19	Constraint migration map	Chart	FRAMEWORK	Author, informed by IEA ^{[22][23]} and IFR ^{[1][2]}
20	Investment opportunity map	Table	FRAMEWORK	Author framework with cited evidence
21	Capital sequencing by evidence horizon	Table	FRAMEWORK	Author framework with cited evidence
22	Company evidence register	Table	Mixed, tagged per row	Company disclosures as cited
23	Global robotics market by segment, 2025A–2030E	Chart	CALCULATION + ILLUSTRATIVE	Market Outlook; IFR unit and market-value data ^{[1][3][52]} , A3 ^{[42][43]} , Intuitive ^[47]
24	Scenario outcomes to 2040 and the changing revenue mix	Chart	CALCULATION + ILLUSTRATIVE	Market Outlook
25	Segment growth to 2040 and share of market expansion	Table	CALCULATION	Market Outlook

Bibliography

All URLs accessed 27–29 August 2026. Reference year of the underlying data is stated separately from the publication date wherever the two differ.

Primary statistical and institutional sources

- [1] International Federation of Robotics. *World Robotics 2025 — Industrial Robots: global robot demand in factories doubles over 10 years*. 25 September 2025. Reference year 2024. <https://ifr.org/ifr-press-releases/news/global-robot-demand-in-factories-doubles-over-10-years>
- [2] International Federation of Robotics. *Robot density surges in Europe, Asia and the Americas*. 2025. Reference year 2024. <https://ifr.org/ifr-press-releases/news/robot-density-surges-in-europe-asia-and-americas>
- [3] International Federation of Robotics. *World Robotics 2025 — Service Robots: service robots see global growth boom*. October 2025. Reference year 2024; 294-supplier sample, not projected to the industry. <https://ifr.org/ifr-press-releases/news/service-robots-see-global-growth-boom>
- [4] International Federation of Robotics. *Global robot density in factories doubled in seven years*. 20 November 2024. Reference year 2023. <https://ifr.org/ifr-press-releases/news/global-robot-density-in-factories-doubled-in-seven-years>
- [5] International Federation of Robotics. *World Robotics 2024 — global market press release*. 24 September 2024. Reference year 2023. https://ifr.org/downloads/press2018/2024-SEP-24_IFR_press_release_World_Robotics_2024_-_global_market.pdf
- [6] International Federation of Robotics. *Record of 4 million robots working in factories worldwide*. 2023. Reference year 2022. <https://ifr.org/ifr-press-releases/news/record-of-4-million-robots-working-in-factories-worldwide>
- [7] International Federation of Robotics. *World Robotics Report: "All-time high" with half a million robots installed in one year*. 2022. Reference year 2021. <https://ifr.org/ifr-press-releases/news/wr-report-all-time-high-with-half-a-million-robots-installed>
- [8] International Federation of Robotics. *Collaborative robots — how robots work alongside humans* (position paper). Reference year 2023. <https://ifr.org/ifr-press-releases/news/how-robots-work-alongside-humans>
- [9] US Bureau of Labor Statistics. *Employer Costs for Employee Compensation — March 2026*. Published 12 June 2026. <https://www.bls.gov/news.release/ecec.nr0.htm>
- [10] US Bureau of Labor Statistics. *Occupational Employment and Wage Statistics, national employment and wage data by occupation, May 2025*. <https://www.bls.gov/news.release/ocwage.t01.htm>
- [11] Federal Reserve Bank of New York, Liberty Street Economics. *The mysterious slowdown in U.S. manufacturing productivity*. July 2024. Data 1987–2022, BLS productivity series. <https://libertystreeteconomics.newyorkfed.org/2024/07/the-mysterious-slowdown-in-u-s-manufacturing-productivity/>
- [19] International Monetary Fund. *Gen-AI: Artificial Intelligence and the Future of Work*. Staff Discussion Note SDN/2024/001, January 2024. <https://www.imf.org/-/media/files/publications/sdn/2024/english/sdna2024001.pdf>
- [20] United Nations, Department of Economic and Social Affairs, Population Division. *World Population Prospects 2024: Summary of Results*. 2024. Medium variant. https://population.un.org/wpp/assets/Files/WPP2024_Summary-of-Results.pdf
- [21] World Bank. *World Development Indicators, series SP.POP.1564.TO (population ages 15–64, total)*. Retrieved via the World Bank API, 29 August 2026. <https://api.worldbank.org/v2/country/CHN;JPN;KOR;DEU;ITA;USA/indicator/SP.POP.1564.TO>
- [22] International Energy Agency. *Rare Earth Elements — executive summary*. 2026. Reference year 2024 for shares; projections to 2035. <https://www.iea.org/reports/rare-earth-elements/executive-summary>
- [52] International Federation of Robotics. *Top 5 Global Robotics Trends 2026*. January 2026. Source of the US\$16.7bn global industrial robot installations market value, reference year 2024. <https://ifr.org/ifr-press-releases/news/top-5-global-robotics-trends-2026>
- [23] International Energy Agency. *Electricity 2026 — executive summary*. February 2026. <https://www.iea.org/reports/electricity-2026/executive-summary>
- [45] China Machinery Industry Federation, as reported by ECNS. *China's robot output in first three quarters surpasses 2024's total*. 6 November 2025. <https://www.ecns.cn/news/cns-wire/2025-11-06/detail-ihewtykv6243775.shtml>
- [46] International Federation of Robotics. *China tops world record of 2 million factory robots*. 25 September 2025. Reference year 2024. https://ifr.org/downloads/press_docs/2025-09-25-IFR_press_release_China_in_English.pdf
- [49] International Labour Organization. *Robots worldwide: the impact of automation on employment and trade*. 27 October 2018. <https://www.ilo.org/publications/robots-worldwide-impact-automation-employment-and-trade>

Peer-reviewed and institutional research

- [12] Acemoglu, D. and Restrepo, P. *Robots and Jobs: Evidence from US Labor Markets*. *Journal of Political Economy*, 128(6), 2020, pp. 2188–2244. NBER Working Paper 23285. <https://www.nber.org/papers/w23285> · DOI 10.1086/705716
- [13] Graetz, G. and Michaels, G. *Robots at Work*. *Review of Economics and Statistics*, 100(5), 2018, pp. 753–768. https://personal.lse.ac.uk/michaels/graetz_michaels_robots.pdf · DOI 10.1162/rest_a_00754
- [14] Dauth, W., Findeisen, S., Suedekum, J. and Woessner, N. *The Adjustment of Labor Markets to Robots*. *Journal of the European Economic Association*, 19(6), 2021, pp. 3104–3153. <https://sebastian-findeisen.com/papers/adjustmentslaborrobots.pdf>
- [15] Aghion, P., Antonin, C., Bunel, S. and Jaravel, X. *What Are the Labor and Product Market Effects of Automation? New Evidence from France*. OFCE Working Paper 2020-01 / CEPR DP14443, 2020. <https://www.ofce.fr/pdf/dtravail/OFCEWP2020-01.pdf>
- [16] OECD. *Determinants and impact of automation*. OECD Social, Employment and Migration Working Papers, 22 February 2019. DOI 10.1787/ef425cb0-en
- [17] OECD. *OECD Employment Outlook 2025 — can we get through the demographic crunch?* 2025. https://www.oecd.org/en/publications/oecd-employment-outlook-2025_194a947b-en.html
- [18] OECD. *Job Creation and Local Economic Development 2024*. 2024. https://www.oecd.org/en/publications/job-creation-and-local-economic-development-2024_83325127-en.html
- [51] OECD. *Labour-saving technologies and employment levels*. 2022. https://www.oecd.org/content/dam/oecd/en/publications/reports/2022/01/labour-saving-technologies-and-employment-levels_c1bb460f/9ce86ca5-en.pdf

Technical and robot-learning publications

- [36] Google DeepMind. *Gemini Robotics 1.5 brings AI agents into the physical world*. 25 September 2025. <https://deepmind.google/blog/gemini-robotics-15-brings-ai-agents-into-the-physical-world/>
- [37] Google DeepMind. *RT-2: new model translates vision and language into action*. 28 July 2023. <https://deepmind.google/blog/rt-2-new-model-translates-vision-and-language-into-action/>
- [38] Open X-Embodiment Collaboration. *Open X-Embodiment: Robotic Learning Datasets and RT-X Models*. 2023. <https://robotics-transformer-x.github.io/paper.pdf>
- [39] Google DeepMind. *Scaling up learning across many different robot types*. 3 October 2023. <https://deepmind.google/blog/scaling-up-learning-across-many-different-robot-types/>
- [40] Physical Intelligence. $\pi_{0.5}$: a Vision-Language-Action Model with Open-World Generalization. 2025. <https://www.pi.website/download/pi05.pdf>
- [41] NVIDIA. *NVIDIA Blackwell-powered Jetson Thor now available, accelerating the age of general robotics*. 25 August 2025. <https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-Blackwell-Powered-Jetson-Thor-Now-Available-Accelerating-the-Age-of-General-Robotics/default.aspx>

Industry statistics and market data

- [24] BloombergNEF. *Lithium-ion battery pack prices fall to \$108 per kilowatt-hour despite rising metal prices*. 2025 Battery Price Survey, 9 December 2025. <https://about.bnef.com/insights/clean-transport/lithium-ion-battery-pack-prices-fall-to-108-per-kilowatt-hour-despite-rising-metal-prices-bloombergnef/>
- [42] Association for Advancing Automation (A3), as reported by The Robot Report. *Q2 2026 robotics demand increased across industries*. 2026. <https://www.therobotreport.com/q2-2026-robotics-demand-increased-across-industries-reports-a3/>
- [43] Association for Advancing Automation (A3). *Robot orders grow 6.6% in 2025 as general industries drive broader automation adoption*. 2026. Reference year 2025. <https://www.automate.org/robotics/news/robot-orders-grow-6-6-in-2025-as-general-industries-drive-broader-automation-adoption>
- [44] Construction Physics. *What progress has there been in industrial robots?* Compiling IFR quality-adjusted and nominal robot price series and historical price records. <https://www.construction-physics.com/p/what-progress-has-there-been-in-industrial>

Company disclosures and primary company communications

- [25] BMW Group. *BMW Group bringing Physical AI to Europe: pilot project at Plant Leipzig; first pilot deployment of humanoid robots successfully completed at Plant Spartanburg, USA*. 2026. <https://www.press.bmwgroup.com/global/article/detail/T0455864EN/>
- [26] BMW Group. *BMW Group advances the use of Physical AI in production with Figure 03 project in Spartanburg*. 2026. <https://www.press.bmwgroup.com/global/article/detail/T0458778EN/>
- [27] Amazon. *Amazon deploys its one millionth robot*. 1 July 2025. <https://www.aboutamazon.com/news/operations/amazon-million-robots-ai-foundation-model>
- [28] Amazon. *Introducing Blue Jay and Project Eluna, Amazon's latest robotics and AI technology for its operations* (including the update of 25 February 2026 stating Blue Jay is no longer in use in operations). <https://www.aboutamazon.com/news/operations/new-robots-amazon-fulfillment-agent-ai>
- [29] Agility Robotics, as reported by The Robot Report. *Humanoid maker Agility Robotics to go public through SPAC merger*. June 2026. Includes company-reported operating hours, bill of materials, capacity and pre-order figures. <https://www.therobotreport.com/humanoid-maker-agility-robotics-go-public-through-spac-merger/>
- [30] Agility Robotics. *Digit moves over 100,000 totes in commercial deployment*. 2025. <https://www.agilityrobotics.com/content/digit-moves-over-100k-totes>
- [31] GXO Logistics. *GXO signs industry-first multi-year agreement with Agility Robotics*. <https://investors.gxo.com/news-releases/news-release-details/gxo-signs-industry-first-multi-year-agreement-agility-robotics>
- [32] Unitree Robotics. *G1 humanoid robot product page*. Accessed 28 August 2026. <https://www.unitree.com/g1/>
- [33] Smart Analytics Global humanoid shipment data and Unitree IPO pricing, as reported by TechTimes and CNBC. *Unitree prices IPO as humanoid leader; data shows AgiBot led H1 shipments*, 10 August 2026; *China's backflipping robot maker Unitree pops 542% in Shanghai debut*, 19 August 2026. <https://www.techtimes.com/articles/323791/20260810/unitree-prices-ipo-humanoid-leader-data-shows-agibot-led-h1-shipments.htm>
- [34] Tesla, Inc. *Annual Report on Form 10-K for the fiscal year ended 31 December 2025*. Filed 28 January 2026. <https://www.sec.gov/cgi-bin/browse-edgar?action=getcompany&CIK=0001318605&type=10-K>
- [35] Electrek. *Tesla pushes Optimus V3 reveal later this year — again*. 22 April 2026. <https://electrek.co/2026/04/22/tesla-optimus-production-fremont-model-sx-line/>
- [47] Intuitive Surgical, Inc. *Fourth quarter and full year 2025 results*. Filed with the SEC, January 2026. <https://www.sec.gov/Archives/edgar/data/1035267/000103526726000006/q425ex-991earningsrelease.htm>

Journalism and secondary reporting (used only where no primary record exists)

- [48] United States Senate Committee on Health, Education, Labor and Pensions (minority). *Sanders demands answers on Amazon's plan to automate warehouses and replace 600,000 workers*. 2025. Reporting New York Times and Reuters accounts of Amazon internal planning documents and a Morgan Stanley savings estimate. <https://www.help.senate.gov/dem/newsroom/press/news-sanders-demands-answers-on-amazons-plan-to-automate-warehouses-replace-600000-workers>
- [50] AgFunderNews. *DailyRobotics gears up for commercial launch in California in 2026 with robotic strawberry harvester*. 2026. All performance figures are company-reported and not independently verified. <https://agfundernews.com/dailyrobotics-gears-up-for-commercial-launch-in-california-in-2026-with-robotic-strawberry-harvester>

DISCLAIMER

This report is provided for information and analysis. It is not investment, legal, tax or accounting advice, and it is not a recommendation to buy, sell or hold any security. The author is not a licensed investment adviser. All models are the author's own and rest on stated assumptions that are not estimates of any specific product or company; they should not be relied on for a commercial decision without independent verification. Company information is drawn from public disclosures and is labelled by evidence status; where a figure is company-reported it has not been independently audited by the author. Readers should consult a qualified professional adviser before acting on anything in this report.