ROBOTICS ECONOMICS

You Cannot Drag and Drop a Robot Into a Job

The $1.66 robot prices machine time. The market buys accepted work.

Tal Cohen | talcohen.ai | August 2026

A widely shared essay, "How to think about the robotics market," makes a powerful case for humanoids. AI is giving machines capabilities that recently belonged to science fiction. Bodies are getting cheaper. Any company paying $20, $30 or $40 an hour for manual labor will feel pressure to buy the same work for less.

I agree with the direction. The economic unit is wrong.

The argument turns on one phrase. A humanoid, it says, will be "drag and dropped" into existing jobs. The robot costs $1.66 an hour because it is assumed to work twenty hours a day, every day, for seven years. The market approaches $60 trillion because one machine is assumed to take the economic place of one worker.

"Drag and drop" is not decoration. It contains the denominator. Delivery gets a robot to the factory gate, and the job waits deeper inside the institution.

What $1.66 actually prices

The division is correct. A $50,000 purchase plus $35,000 for maintenance and electricity, divided by 51,100 assumed working hours, yields $1.66.

But that is a cost per assumed working hour, not a cost per accepted unit of work. The calculation has already granted availability, a steady supply of tasks, throughput, quality and permission to operate.

Hold its $85,000 lifetime cost constant. If half the assumed capacity becomes accepted work, the number doubles to $3.33. At one quarter, it becomes $6.65. At ten percent, it becomes $16.63. That is before financing, integration, software, supervision, tooling, insurance, rework and uninsured loss. Some maintenance and energy would fall with use, so these are sensitivities, not forecasts. The denominator can still move the result by an order of magnitude.

That may still beat human labor. It is not $1.66.

Industrial robotics has priced the gap between a machine and a working system for forty years. The industry's own guidance, citing Boston Consulting Group, budgets at least three times the robot's price to reach production.

An hour is also a poor unit when outputs differ. A robot that moves sixty totes within specification is not equivalent to one that moves thirty and needs four interventions, even if both run for an hour. The honest comparison is fully burdened cost per accepted unit of equivalent output.

The word accepted matters. A completed movement becomes economic output only when it meets the operating standard, enters the production record and is allowed to continue.

What production looks like

BMW's Figure 02 deployment is useful because it is real. Over an eleven-month program, the deployment moved more than 90,000 components, logged about 1,250 operating hours and supported production of more than 30,000 vehicles. The disclosures establish meaningful production. They do not provide enough per-robot runtime, downtime or accepted-output data to calculate steady-state utilization.

They do fix the schedule. BMW ran ten-hour shifts, Monday through Friday. Take that documented 50-hour week as a sensitivity: 18,200 scheduled hours over seven years, against the 51,100 the thesis assumes. Hold the $85,000 lifetime cost fixed and $1.66 becomes $4.67. Still potentially compelling, and nearly three times the headline, before downtime, interventions, rejected work, integration and support. Scheduled hours are not accepted work either.

Note what the reported runtime is close to. Six months of daily ten-hour shifts is roughly 1,300 scheduled hours, and Figure reports more than 1,250. Across that stretch the machine was working close to all the time it was allowed. The schedule bound the output, not the robot.

BMW also names what stood behind the result: production IT, occupational safety, process management, shop-floor logistics and standardized interfaces. The body shop was chosen partly because it already had deep automation experience.

Nobody dragged and dropped that robot into a job. BMW and Figure built the conditions under which it could work.

Utilization is not a robot specification. It is a system outcome.

The missing economic unit

Choose the unit of accepted output first: a tote moved within specification, a weld that passes inspection, a task-hour of defined quality. Then watch one robot for one shift. Its spec sheet allows 600 totes. It physically completes 480, because grippers jam and an update forces a reboot. The institution accepts 400, because twenty landed outside the marked zone and sixty lack a clean safety record. The day obeys one identity:

Accepted output = nominal capacity × execution yield × Habitat yield

Execution yield is 480 of 600. Habitat yield is 400 of 480. These are sequential accounting ratios, not independent causal shares. The split shows where the output disappears, and the reasons behind each loss show who can act. Both the robot and the Habitat can move either ratio.

Then price it honestly. Put the full yearly bill on top: the annual capital charge for the robot and deployment assets spread over their useful life, plus software, supervision, insurance and rework. Put the year's accepted output underneath:

Cost per accepted unit = annualized total system cost ÷ annual accepted output

Total system cost annualizes robot and deployment capital over its useful life, then adds recurring software and human support, maintenance, insurance premiums, rework and uninsured expected loss. The comparison must be symmetric. Do not charge the robot for the whole factory while treating human management and safety as free. Compare incremental costs on both sides, or fully burdened system costs on both sides.

The human system is mature and mostly inherited. A worker arrives inside language, training, management, safety practice, insurance, law and settled responsibility. A robot needs a machine-specific extension of that system, built at this site, in this quarter, on the buyer's money.

Economists call such an investment complementary capital. Electricity needed factories rebuilt around motors. Computers needed workflows and skills rebuilt around software. Physical AI needs an additional complement because it does not merely advise. It acts.

I call that complement the Operating Habitat: the persistent organizational, technical and governance environment that makes physical AI deployable, governable and improvable even when models, machines or vendors change. In economic terms, it converts machine capability into accepted work. It defines the work and its limits, grants authority, preserves an independent record and assigns consequences.

Habitat enters the numerator because it costs money to build and operate. It also changes the denominator by determining how much technical capacity becomes accepted work. The drag-and-drop model sets Habitat cost to zero and Habitat yield to 100 percent.

Risk enters twice. Routine failure and rework add expected cost. Risk beyond the institution's tolerance closes the gate. A robot that cannot be authorized produces no economically usable output.

What this model predicts

First, it predicts which bodies win.

A humanoid is a bet that changing the machine will cost less than changing the world. Human form can reuse doors, aisles, stairs, tools and workstations built for us. That advantage is strongest in variable brownfield environments. In a new plant built around stable work, fixed arms, conveyors, mobile robots or wheeled manipulators may deliver more accepted output with less complexity.

The humanoid bet has a second leg, and it is the stronger one. The same machine can learn new work. A demonstration replaces a reprogramming project, the fleet shares what one robot learns, and next year the same hardware does more work than it does today. BMW saw a version of this: motions trained in the laboratory reached stable shift work faster than expected. Figure has since retired the Figure 02 fleet and carried those lessons into its successor. Learning is the strongest force pushing real output toward the assumed hours. It raises execution yield on capital already paid for, and it spreads one Habitat across more tasks.

Learning also changes behavior, and behavior is what the institution accepted. To the safety case and the insurer, a robot that behaves differently after an update is a different robot. The machine retrains fast. The institution then re-accepts: it verifies the new behavior, confirms the bounds still hold, and brings the record current. The strongest system is a humanoid that learns quickly inside a Habitat that keeps up.

A humanoid can fit the building and still fail to fit the institution. Generality earns its premium only when broader task coverage and redeployment create enough utilization to pay for the complexity. The winner is the system that produces the most risk-adjusted value from accepted output.

Second, the model changes what scale means.

The first deployment absorbs task design, interfaces, validation, evidence, training and operating rules. Delivery volume is manufacturing scale. Deployment scale begins when some of that investment survives and the next comparable deployment is cheaper and faster.

This changes procurement. A $50,000 robot that takes nine months to reach accepted production can cost more than an $80,000 robot admitted in six weeks. Time consumes financing, management attention and useful economic life. The best robot may be the one that reaches accepted production fastest and forces the institution to rebuild the least.

Third, the model predicts who captures value.

Robot makers win when their interfaces, failure behavior, records and service reduce the customer's deployment burden. Large operators win when they reuse the same operating system across tasks and sites. Integrators build a compounding advantage only when bespoke projects become repeatable infrastructure.

The durable advantage may belong to whoever owns the layer that improves with every deployment and survives a model or body swap.

What $60 trillion measures

The global labor pool describes immense pressure to automate. It is not a robotics revenue forecast.

The $60 trillion figure is an annual flow of labor income. The $200 billion set against it is a stock of private-company value. Dividing dollars per year by dollars does not produce a meaningful “300x” market multiple.

The two sides do not describe the same machines either. The essay’s own footnote assembles the $200 billion from robotics and autonomy companies of every kind, from robotaxis to defense, with humanoid makers a modest slice.

Labor income also maps poorly onto robot revenue. Jobs are bundles of tasks, many cognitive or hybrid. A machine working several shifts is not equivalent to one employed person. If automation cuts the price of a task, the wage pool does not transfer intact to suppliers. Some value goes to customers and consumers. Lower prices may create more demand, but greater output is not the same as captured revenue.

Treat $60 trillion as a pressure field. It explains why the prize for solving deployment is large. The realized supplier market is simpler: contracted price per accepted unit multiplied by accepted units, plus any fixed hardware or service fees.

Leasing makes the correction impossible to avoid. Once a provider sells labor rather than hardware, the invoice needs a unit: an accepted hour, completed task, throughput target or uptime commitment. Someone must meter it, verify it and carry the failure risk. Seven years of invoices are not seven years of profit. A fixed lease shifts financing and residual-value risk toward the fleet owner, while utilization, uptime and maintenance risk follow the contract. Labor-as-a-service and performance pricing move more of those toward the provider.

The record becomes part of the product.

The stronger robotics case does not need to call $1.66 a fully burdened wage or $60 trillion supplier revenue. It needs a measurable passage from capacity to accepted output, and evidence that each deployment makes the next one easier.

AI gives the robot a brain. Manufacturing gives it a body. Neither gives it a job.

A job exists when an institution can define the work, admit the machine, accept the output and answer for what happens.

A better model, you can buy. A Habitat, you have to build.

Sources

Original essay. intern, “How to think about the robotics market,” X, August 2026. https://x.com/intern/article/2089418795353796754

Cost benchmark. Association for Advancing Automation, “Calculating robot ROI: How to determine the true cost of robotics,” citing Boston Consulting Group. https://www.automate.org/robotics/editorials/calculating-robot-roi-how-to-determine-the-true-cost-of-robotics

Production case. BMW Group, “BMW Group to deploy humanoid robots in production in Germany for the first time,” February 27, 2026. https://www.press.bmwgroup.com/global/article/detail/T0455864EN/bmw-group-to-deploy-humanoid-robots-in-production-in-germany-for-the-first-time?language=en

Economic precedent. Erik Brynjolfsson, Daniel Rock and Chad Syverson, “The Productivity J-Curve,” NBER Working Paper 25148, revised January 2020. https://www.nber.org/papers/w25148