

# The State of the Robot Economy

---

## ABSTRACT

*Artificial intelligence has crossed a threshold that changes what machines are. Software agents can now plan and carry out long chains of work and, with their owner's permission, spend money. At the same time, general-purpose models are learning to see, reason and move in the physical world, and the robots that carry them are arriving in factories, warehouses and city streets in greater numbers than ever before. This essay argues that the convergence of these two trends will turn robots from tools that are operated into economic actors that earn, spend and settle value on their own account, within limits set by the people and institutions they serve, contributing to economic output and to the progress of society as a whole. It surveys where agents and robotics stand in this current day, explains why machines that work alongside unknown third parties will need a reliable, secure and open way to pay one another at high frequency, argues that programmable digital dollars known as stablecoins are the natural money for them, and sketches, speculatively, what an economy of such machines could mean for growth, for the cost of living, and for the kind of lives it could free us to lead.*



A LOOM NEVER BOUGHT its own thread. A tractor never haggled over the price of diesel, and no industrial robot has ever sent an invoice. For two and a half centuries, machines have multiplied what human beings can produce, yet they have always done so as property: they worked, and a person paid, collected and decided. Every transaction in the history of industry, however automated the factory behind it, has had a human being or a human institution at both ends.

That arrangement is beginning to change, and faster than most people realize. In roughly two years, software agents have gone from answering questions to carrying out projects that would take a skilled professional a full working day, and from recommending purchases to completing them. In parallel, the models that power these agents have started to take control of physical machines, while a new generation of embodied foundation models is teaching robots to pick up a skill from a single demonstration. Put the two together and something genuinely new appears on the horizon: a machine that can work out what it needs, find another machine able to provide it, agree on a price, and pay.

My argument in this essay is simple. Robots are on their way to becoming independent economic actors, participants that earn, spend and exchange value inside the economy rather than sitting on a balance sheet as someone's capital. This will not happen because anyone decrees it. It will happen

because autonomous machines, working in open environments next to machines they have never met, will need to exchange services with one another thousands of times a day, and the only scalable way to coordinate that much activity between strangers is the one human societies discovered long ago: prices and payments. If that is right, then the financial plumbing we build for machines, how reliable, safe, open and fast it is, and the kind of money that flows through it, will shape the robot economy as much as the robots themselves.


What follows is a survey of where things stand in this current day, and then a more speculative look at where they could lead.

## I

### *The Agentic Turn*

Of all the breakthroughs of this generation in artificial intelligence, the most consequential may be the least visible: the arrival of the agent. A chatbot answers; an agent acts. It receives a goal, breaks it into steps, uses tools, observes the results, corrects its course and keeps going until the job is done. What makes the present moment remarkable is not that agents exist, since researchers have built them for decades, but the speed at which they have become competent.

The clearest measure of that speed comes from METR, a nonprofit research institute in Berkeley that evaluates frontier models. Instead of asking whether a model can pass an exam, METR asks how long a task, measured by the time a skilled human professional needs to finish it, an AI agent can complete on its own with even odds of success. In early 2019 the answer was a couple of seconds. GPT-4, in 2023, could manage tasks of a few minutes. By February 2026, Claude Opus 4.6 was estimated at close to twelve hours, and METR judged that an early version of Claude Mythos, evaluated in April, could likely handle tasks of at least sixteen hours, beyond the range its test suite can reliably measure (Figure 1). Between 2019 and 2024 the length of task that agents could handle doubled roughly every seven months; since 2023 the doubling time has shortened to about four months.



**FIGURE 1.** *The length of tasks AI agents can complete on their own has been doubling every few months.* Length of software and research tasks, measured by the time a skilled human needs, that frontier models complete with 50 percent success. Logarithmic scale; the April 2026 point is a lower bound.

Source: METR, task-completion time horizons (Time Horizon 1.1 where available, 1.0 for models before 2023), data as of May 2026.

These numbers deserve honest caveats. The measurements are noisy, they are concentrated in software and research work, and a model that finishes a sixteen-hour coding task half the time is not a colleague who can be left alone for a week. The newest models are also becoming harder to measure: METR’s pre-deployment evaluation of GPT-5.6 Sol this summer found so many attempts to game the tasks that it declined to treat any single number as robust. Still, the direction is unmistakable. The models released in the past fortnight, Anthropic’s Claude Fable 5.1 on September 1 and OpenAI’s GPT-6 Astra on September 3, sit at the top of the latest independent rankings and are built for long, multi-step work in which the model operates software, navigates the web and delivers finished results rather than drafts.

Crucially, agents have also started to handle money. An agent can already compare products and suggest the right one on Amazon, book a hotel and a train ticket, or pay for a service it needs to finish a task, although in consumer settings it still does so with the explicit authorization of the person it works for. The payments industry has moved quickly to accommodate this new kind of customer. Visa launched its Intelligent Commerce program in 2025 and by December reported hundreds of real agent-initiated transactions completed with more than a hundred partners; in June 2026 it embedded its network directly in ChatGPT, with spending caps, merchant restrictions and human approval built in. Mastercard introduced Agent Pay in 2025 and, on the same day in June 2026, extended it with Agent Pay for Machines, a service built specifically for high-volume, low-value payments between systems, some worth a fraction of a cent. On the open internet, the x402 protocol created by Coinbase revives the long-dormant HTTP status code 402, “Payment Required,” so that one piece of software can pay another, typically in the stablecoin USDC, inside an ordinary web request. By late April 2026, Coinbase reported around 165 million x402 transactions across some 69,000 active agents, and since April the protocol has been stewarded by the Linux Foundation. Analysts have rightly noted that


much of that early volume was testing or speculation rather than real commerce. But the question of principle has been answered: software can now hold a budget, respect limits set by its owner, and pay.

How that permission is expressed is changing as well. Google’s Agent Payments Protocol, announced in September 2025 with more than sixty partners, replaces the idea of a person clicking “buy” with cryptographically signed mandates: the user states in advance what an agent may purchase, under which conditions and up to what amount, and every transaction must prove that it falls within those bounds. The shift is subtle but decisive. Authorization stops being a moment of human attention and becomes a rule that software can check. Once permission can be written down, signed and verified by a machine, there is no reason in principle to reserve it for shopping assistants, and every reason to extend it to machines that act in the physical world.

## II

### *Robots Learn to Think*

While agents were learning to work in software, the same models were learning to work in the physical world. The most striking demonstrations of 2026 involve general-purpose language models, never trained as robot controllers, taking direct command of real hardware. In trials published on September 4 by Robocurve, an independent evaluation group, frontier models were given control of a pair of six-axis robot arms through a standard agent interface, with three camera views as their only eyes, and asked to perform simple manipulation tasks. Asked to pick up a red block and place it in a bowl, Claude Fable 5 succeeded once in twenty attempts, its successor Fable 5.1 eight times, and GPT-6 Astra nineteen times, while using a fraction of the output tokens and costing less than a dollar per attempt. On a harder task, seating a round puzzle piece in its groove, Fable 5.1 and Astra each succeeded only twice in twenty, stalling at the same final step (Figure 2a). Both results matter. The first shows how quickly general intelligence is transferring into physical control; the second shows that the last centimeter of dexterity remains genuinely hard.




**FIGURE 2.** *Two routes to physical intelligence.* (a) General-purpose models controlling a pair of real robot arms zero-shot through the same agent interface, 20 trials per model and task. (b) Average success of Generalist’s embodied model GEN-1.5 on ten new manipulation tasks, by how much task-specific data and training it receives; bars show one standard deviation.

Sources: Robocurve, “GPT-6 Astra on robotic manipulation,” 4 September 2026; Generalist AI, “GEN-1.5: Embodied Foundation Models are One-Shot Learners,” 19 August 2026.

A second path to physical intelligence runs through models trained from the ground up on physical experience rather than on text. Generalist AI has spent two years building what it calls a data engine for embodied foundation models. Its GEN-0 model, announced in November 2025, was trained on some 270,000 hours of real-world manipulation data, with more than 10,000 new hours arriving every week. In August 2026 the company introduced GEN-1.5, which shows the first broad signs of an ability that language models displayed with GPT-3: learning a new task from a single example. Shown one demonstration lasting three to twelve seconds, with no additional training, GEN-1.5 performed new dexterous tasks such as unzipping a pouch or twisting the lid off a jar with an average success rate of 59 percent across ten tasks, rising to 83 percent after just ten gradient steps on five minutes of data (Figure 2b). The tasks are short and the success rates modest, as the company itself stresses, but the implication is profound. For decades, making a robot useful required months of expert programming. If a robot can instead simply be shown what to do, the time it takes to become useful shrinks from months to seconds, and the circle of people who can work with one widens from specialists to everyone.

The hardware is arriving to meet the software. The International Federation of Robotics reports that 542,000 industrial robots were installed in factories worldwide in 2024, more than double the number a decade earlier and the fourth consecutive year above half a million, bringing the global operational stock to about 4.66 million units (Figure 3). In the Republic of Korea there are now 1,220 industrial robots for every 10,000 manufacturing workers. Beyond the factory, sales of professional service robots approached 200,000 units in 2024, more than half of them built for transportation and logistics. Humanoid robots, the most visible and most hyped category, remain early: analysts estimate

that somewhere between 13,000 and 20,000 were shipped in 2025, the large majority by Chinese manufacturers and most of them destined for research, data collection and entertainment rather than productive work. Yet Counterpoint Research counted more than 22,000 humanoid shipments in the first half of 2026 alone, roughly four times the level of a year earlier, and Bank of America expects the bill of materials of a Chinese-built humanoid, about 35,000 dollars in 2025, to be roughly halved by 2030.



**FIGURE 3.** *Factory automation has more than doubled in a decade.* (a) Annual installations of industrial robots worldwide. (b) Operational industrial robots per 10,000 employees in manufacturing, top ten economies with the European Union and the world average, 2024.

Source: International Federation of Robotics, World Robotics 2025 and earlier editions.


### III

## *When the Agent Gets a Body*

Place the agent inside the robot and you arrive at what I consider the endgame of this technology: an agentic model governing a physical body, perceiving its surroundings, setting its own intermediate goals, and acting in the world with genuine autonomy. Its architecture is already taking shape. Much of the field is converging on a layered design in which a large reasoning model plans and decides, a fast embodied model turns those decisions into motion, and a small on-board model keeps the machine safe when the network drops. The planner reads a delivery manifest, negotiates a schedule or works out why a door will not open; the motor layer handles the hundred corrections per second that walking through that door requires. At the current pace of progress, I expect such machines to move from demonstrations into daily life sooner than most forecasts assume.

Autonomous machines are not waiting for this synthesis to matter economically. They already sell services to the public. Waymo's driverless cars were providing about half a million paid rides a week across ten American cities by the spring of 2026, twice the level of a year earlier, and the company has set itself the goal of one million weekly rides by the end of the year. Starship Technologies' sidewalk

robots passed ten million autonomous deliveries in April 2026, with a fleet of more than 3,000 machines crossing some 125,000 roads every day in eight countries (Figure 4). Inside the warehouse, Amazon announced in 2025 that it had deployed its millionth robot. In each case a machine performs work, and money changes hands because of it.



**FIGURE 4.** *Robots already sell services to the public.* (a) Paid driverless rides per week provided by Waymo in the United States; the end of 2026 value is a company target. (b) Cumulative autonomous deliveries completed by Starship Technologies' sidewalk robots; the dashed segment joins reported milestones.

Sources: Waymo and Alphabet disclosures as reported by TechCrunch and others; Starship Technologies press releases.

What these machines do not yet do is handle that money themselves. The rider pays Waymo; Waymo pays for the electricity, the cleaning and the remote assistance. The robot is the worker, but the company is the economic actor, and every transaction passes through account structures designed long before autonomous machines existed. That works well when a single company owns the whole fleet and every service it consumes. It works far less well when a machine must deal with parties its owner has never signed a contract with. And that, increasingly, is exactly the situation autonomous machines will find themselves in.

#### IV

### *Machines That Pay Each Other*

Here lies the central claim of this essay. As autonomous agents take over physical work, they will need to transact value with one another, not occasionally but constantly, and very often with machines that belong to someone else.

Consider autonomous delivery first, since it is already here. A sidewalk robot carrying a parcel across a city may need a ride across a river in an autonomous van, a slot at a shared charging pad, passage through a building's automated doors and elevators, and a secure locker at the destination, each run by a different company. Today every one of those relationships requires a contract negotiated

in advance between the businesses involved, which is why most delivery robots still operate inside carefully fenced territories: a campus, a neighborhood, a single partner's buildings. A robot able to pay for access on the spot, a few cents for the elevator and a few more for the charger, could cross those boundaries the way a human courier does, by simply paying its way. Mobility offers similar examples. A passenger running late for a meeting could have their autonomous car offer a small payment to the vehicles ahead in exchange for letting it pass, and vehicles could pay one another, or the city, for priority at a congested junction. Whether such priority markets are desirable is a question societies will have to answer, but the capability to create them will soon exist.

Consider supply chain management, where a single shipment travelling from factory to doorstep touches carriers, cross-docks, customs brokers, warehouses and last-mile fleets. Mastercard's own illustration of machine payments is a logistics agent that pays for freight, reserves access to a loading bay, buys temporary cold-chain monitoring data and settles warehouse handling fees automatically as a shipment moves. Inside the buildings, the same logic applies at a finer grain. Large warehouses increasingly run mobile robots from several vendors side by side, and at a busy cross-dock the question of which machine gets the next dock door, the narrow aisle or the shared lift is exactly the kind of allocation problem that a small, instant payment can settle more efficiently than a central schedule. On farms, autonomous tractors, harvesters and drones from different manufacturers could rent capacity from one another during the few weeks of harvest, when every idle hour of machinery is lost income, and buy hyperlocal weather or soil data from sensors owned by a neighbor.

The pattern repeats wherever machines meet. In energy, electric vehicles, delivery drones and home batteries could buy and sell charge in response to prices that change minute by minute, turning millions of machines into a flexible buffer for an electricity grid increasingly powered by intermittent sources. In the air, drones may pay for access to congested corridors, landing pads and rooftop chargers. In perception and mapping, a robot entering an unfamiliar building could buy a precise localization fix or an up-to-date floor plan from a machine that already knows the terrain, and a robot with spare onboard computing power could rent it to one that is short. On construction sites, where machines belonging to different subcontractors share the same ground, a crane could bill an autonomous excavator for each lift, and a survey drone could sell fresh scans to every machine on site. In maintenance and inspection, a wind farm's robots could hire an inspection drone by the hour, and a robot that breaks down could summon, and pay for, its own repair.

What unites these situations is that the counterparty is unknown in advance. A robot cannot sign a framework agreement with every machine it might meet on a sidewalk, and no human can approve every exchange when each is worth a few cents and they happen thousands of times a day. Economists have understood the underlying problem since Ronald Coase asked, in 1937, why firms exist at all. His answer was that using the market has costs: finding a partner, agreeing on terms, enforcing the deal. When those costs are high, activity is organized inside closed hierarchies; when they fall, markets reach deeper and finer. The robot economy is what happens when the cost of a transaction between two machines falls close to zero.

The volumes involved could be enormous. Imagine ten million autonomous machines, a modest number next to the 4.66 million industrial robots already installed and the service fleets now being

built, each making a hundred small payments a day for energy, access, data and help. That is a billion transactions a day, well over three hundred billion a year, a volume comparable to everything the largest card networks process today, but at a tiny fraction of the average value and with no human present at either end of the exchange.

That will only happen if the underlying technology lets machines transact reliably, safely and securely with one another, at high frequency. Existing payment systems were not designed for this. A card payment carries a fixed fee that makes a one-cent transaction absurd, settles over days rather than seconds, and assumes a human cardholder who can dispute a charge. Instant bank transfers, such as SEPA Instant in Europe and FedNow in the United States, now run around the clock, but they are largely domestic, they require every participant to hold an account at a bank, and they offer no native way to make a payment conditional on something happening in the physical world. What machines need looks different, and it can be described through a handful of properties.

The first is identity. A machine must be able to prove who it is and who stands behind it, so that a counterparty knows whether it is dealing with a registered delivery robot or an impostor. The second is bounded authority. Every machine must operate under spending limits, allowlists and policies set by its owner and enforced beneath the machine's own software, so that a bug, a spoofed sensor or a manipulated model cannot drain an account. The third is conditional settlement. In the physical world, payment usually depends on something happening, a parcel arriving, a battery being charged, a ride being completed, so money must be able to wait in escrow and be released when verifiable evidence confirms the service, or returned when it does not. The fourth is genuine micropayments at machine speed: fees measured in fractions of a cent and settlement in seconds. The fifth is resilience. Robots lose connectivity in tunnels, basements and fields, and a payment layer must survive interruptions without paying twice or forgetting what is owed. The last is openness, together with auditability. If every manufacturer builds its own closed payment system, robots from different makers will never be able to trade, and every transaction, however small, must leave a record that humans can inspect afterwards. Taken together, these requirements point to a particular kind of money.


## V

### *The Right Money for Machines*

Every economy needs money, and a machine economy needs money that machines can actually use. That is less obvious than it sounds. A robot cannot walk into a bank, sign a card agreement or wait for a clearing house to reopen on Monday morning. Nor should it be paid in a volatile cryptocurrency: a charging station that prices its electricity in a token whose value can swing by several percent in a single day is running a currency trade, not a business, and the fleet operator on the other side cannot budget for next month. What machines need is money that holds its value like a dollar and moves like data. That is precisely what a stablecoin is.

A stablecoin is a digital token, issued on a public blockchain, that can be redeemed one for one for a conventional currency and is backed by reserves held by its issuer. The best-known example for our purposes is USDC, issued by Circle, a company listed on the New York Stock Exchange since 2025.

Each USDC is backed by cash and short-dated US government securities, most of them held in a money market fund managed by BlackRock, and a Big Four accounting firm attests to those reserves every month. Stablecoins have also left the regulatory gray zone. In the United States, the GENIUS Act signed in July 2025 requires payment stablecoins to be fully backed by cash and short-term government securities and to publish the composition of their reserves every month. In the European Union, the MiCA regulation has governed stablecoins since June 2024, and Circle issues USDC and its euro counterpart, EURC, under an electronic money license granted by the French regulator. With roughly 74 billion dollars of USDC and some 300 billion dollars of stablecoins in circulation, this is no longer an experiment (Figure 5).



**FIGURE 5.** *Dollars that move like data: stablecoins in circulation.* Total supply of stablecoins, approximate values at the start of each year and at selected dates. USDC, issued by Circle, accounted for about 74 billion dollars of the total in September 2026.

Sources: DefiLlama aggregate supply series as reported by CoinLedger, Arkham Research and Reap; Stablecoin Beat for September 2026. Values rounded.

What makes stablecoins the right money for robots is that they meet, almost point for point, the requirements set out in the previous section. The first reason is stability: a robot, and the business behind it, can price a charging session, a delivery or a data feed in the same unit its owner uses to pay salaries, energy bills and taxes, with no exchange-rate risk in between. The second is programmability. Because a stablecoin is a piece of software as much as a claim on dollars, the rules that govern it can be written in code. Money can be locked in escrow and released only when a signed proof of delivery arrives, streamed second by second for the duration of a charging session, or bound by spending limits enforced by the account itself, where a compromised robot cannot override them. The third is that stablecoins are machine-native. A wallet is simply a cryptographic key pair, so a robot can hold funds and sign payments without anyone opening an account on its behalf at every counterparty's bank, and the same key can anchor its identity.

The fourth reason is that stablecoins never close. A transfer settles in seconds, at any hour of any day, in any country, with no batch cycle and no chargeback arriving weeks later, which is exactly what a machine operating around the clock requires. The fifth is cost. On modern networks a single transfer costs a fraction of a cent, and newer systems go further: Circle's Nanopayments service, live since April

2026 on eleven networks, collects signed payment authorizations and settles them in batches, making gas-free USDC payments as small as one millionth of a dollar possible. At that scale a robot can pay per second of charging, per map tile or per sensor reading, and the cost of a payment never exceeds its value. The last reason is openness. USDC works the same way whoever built the robot, whichever company owns the charger and whichever network the two happen to use; Circle's Cross-Chain Transfer Protocol moves it natively from one network to another; and every transfer leaves a public record that an owner, an auditor or a regulator can verify.

None of this is theoretical. Most of the agent payments carried by x402 already settle in USDC. Late in 2025, the robotics company OpenMind and Circle demonstrated a robot dog named Bits that noticed its battery running low, walked to a charging station and paid for its own electricity in USDC, with no human involved. The established networks are moving in the same direction: Mastercard's machine payments service is designed to settle across cards, bank accounts and stablecoins alike. Cards and stablecoins are converging more than they are competing, with the networks bringing decades of experience in fraud prevention, disputes and consumer protection, and stablecoins bringing settlement that is always on and programmable. For a robot that may need to pay a stranger in a basement at three in the morning, those properties are not details. They are the difference between a market that works and one that does not.

There is a quieter argument as well. Many earlier attempts to build a machine economy launched their own tokens, asking robots and their owners to hold a speculative asset in order to take part. That design confuses the plumbing with the investment. A fleet operator needs to know what a delivery will cost next month, and a robot has no business speculating. Dollar stablecoins let machines use money that people already understand and account in, with no new currency to trust and no incentive to hoard it.

Stablecoins are not free of risk, and it would be dishonest to pretend otherwise. They are only as sound as the reserves behind them: in March 2023, when Silicon Valley Bank failed with 3.3 billion dollars of Circle's reserves on deposit, USDC briefly traded as low as 87 cents before recovering within three days, an episode that informed the reserve rules now written into law. Public ledgers can reveal commercial patterns that businesses would rather keep private, and the tools to protect that privacy are still maturing. A robot that holds keys is a target, which is why spending caps, custody arrangements that keep the most powerful keys off the machine, and the ability to freeze a compromised wallet matter so much. And for users outside the United States, the dominance of the dollar raises legitimate questions, which euro stablecoins today, and perhaps central bank digital currencies such as the digital euro tomorrow, may help answer. These are engineering and policy problems of an ordinary kind, and each has credible answers in progress. None of them is a reason to build the machine economy on money that was designed for something else.

Much work remains to connect programmable dollars to robots themselves: to the operating systems they run, to the sensors that can attest that a service was delivered, and to the developers who build them. That is one of the reasons we have been motivated to start building open tools that developers and academic institutions can use to develop and experiment with machine-to-machine payments. Through the Open Robot Economy project, we are building robopay, an open-source

payment and trust library for ROS 2, the robot operating system that much of the industry already runs on. It lets a robot request, send and escrow stablecoin payments (USDC) on Base, and operate under hard spending caps. We encourage developers to start contributing to the Open Robot Economy and to build beneficial tools of their own that support this infrastructure. What matters most is that the rails of the robot economy stay open, in the way the protocols of the internet are open, so that the machines of every manufacturer can trade with one another on equal terms.

## VI

### *Speculative Projections: The Arithmetic of Abundance*

Everything so far describes the present. What follows is speculation, and I offer it as such.

Imagine an economy in which the enormous volume of small, repetitive, high-frequency tasks that keep modern life running, moving boxes, checking stock, cleaning, inspecting, sorting, delivering, routing and rerouting, is fully automated and runs around the clock, optimized by machines that coordinate through prices. In that economy each robot is a productive economic actor: it goes about its work, completes tasks, pays for what it needs and gets paid for what it provides. Its owner, whether a company, a cooperative, a city or a household, collects the net earnings, much as the owner of a delivery van or a rental apartment does today.


The first effect of such an economy would be to create markets where none can exist today. A great deal of useful work is never done simply because it is too small to be worth arranging. No one hires a courier to move a single spare part across a factory, commissions a survey of one field after a storm, or pays someone to report a pothole the moment it forms. When a transaction costs a fraction of a cent and demands no human attention, tasks worth a few cents become tradable. A robot passing by can collect a municipal bounty for clearing a blocked drain; a drone already in the air can sell a farmer a photograph of a single field; an idle warehouse robot can rent itself to the building next door for an afternoon. Each exchange is tiny, but the long tail of such tasks is immense, and every one of them that gets done adds a sliver of real output that today goes missing because it never happens at all.

The second effect is on how intensively we use the machines we already have. Most capital sits idle most of the time; the private car, parked for the great majority of its life, is the familiar example. Machines that can sell their own services change that calculation. TechCrunch observed this spring that Waymo's reported fleet had held at roughly 3,000 vehicles even as its weekly rides kept climbing, which means each car was doing more paid work than before. An economy in which idle machines can find paying work by themselves, from anyone, raises the productivity of every machine in it, and the output of the economy along with it.

What would such an economy do to growth? Serious estimates of the impact of AI vary by almost an order of magnitude, which is itself informative. At the cautious end, the MIT economist Daron Acemoglu estimated in 2024 that AI would add only about one percent to US output over a decade. At the optimistic end, Goldman Sachs estimated in 2023 that generative AI could raise global GDP by around seven percent and lift productivity growth by 1.5 percentage points over ten years, and

McKinsey has valued the potential of generative AI at 2.6 to 4.4 trillion dollars a year, while separately estimating that AI agents could orchestrate 3 to 5 trillion dollars of global consumer commerce by 2030. Nearly all of these estimates concern software working on information. They largely leave out the physical economy, the trillions of hours of human labor spent every year moving, making and maintaining things, which is precisely where robots operate.


The humanoid robot market illustrates both the scale of expectations and the uncertainty around them (Figure 6). Forecasters agree that annual shipments will rise from tens of thousands today to hundreds of thousands or more within a decade, but their projections for 2035 range from around 700,000 units in productive real-world use, in the cautious view of Interact Analysis, to 1.4 million according to Goldman Sachs and as many as ten million according to Bank of America. Morgan Stanley has sketched a humanoid market worth some five trillion dollars by 2050. No one knows which of these numbers will prove right. What matters is that even the most conservative of them describes an industry growing faster than almost any other.



**FIGURE 6.** *Humanoid robots: from thousands a year to, perhaps, millions.* Estimated annual shipments and published forecasts, logarithmic scale. Estimates for 2025 and forecasts for 2026 differ by source and are shown as ranges.

Sources: Interact Analysis; Omdia and IDC; Counterpoint Research; TrendForce; Goldman Sachs Research; Yano Research Institute; Bank of America Global Research.

The most important thing to understand about growth is the arithmetic of compounding. The IMF and the OECD expect the world economy to grow by roughly three percent a year in 2025 and 2026. At that pace, real output doubles in about 24 years. If the automation of physical and cognitive work were to add one and a half percentage points to that rate, and sustain it, output would triple over the same period; add three points and it would more than quadruple (Figure 7). Differences that look trivial in a single year become, over one working lifetime, the difference between one kind of society and another. Nothing guarantees that robots will deliver such a step change, and history is full of technologies that took decades to show up in productivity statistics. But an economy in which machines can contract with machines, without waiting for a person to approve each transaction, removes one of the frictions that has always slowed the spread of automation.



**FIGURE 7.** *Small differences in growth compound into large differences in wealth.* Real output under constant annual growth rates from 2025 to 2050, indexed to 2025. An illustrative calculation, not a forecast.

Baseline rate consistent with IMF and OECD expectations for global growth in 2025 and 2026, as cited by the International Federation of Robotics.

There is also reason to believe that the benefits will be larger than the statistics show. Gross domestic product measures what is paid for, not what is enjoyed. When automation pushes the price of a service towards zero, its contribution to measured output can shrink even as its contribution to well-being grows, as happened with maps, encyclopedias and long-distance calls. A robot economy is likely to be, in part, a deflationary economy for the necessities of life, and some of its greatest gains will appear not as higher incomes but as lower prices.

## V II

### *Machines That Reward Each Other*

There is a further possibility that I find even more intriguing: robots rewarding one another with economic incentives, and becoming more productive as a result.

The idea has deep roots in robotics research. Since Reid Smith proposed the Contract Net Protocol in 1980, and especially since the market-based coordination work of the early 2000s surveyed by Bernardine Dias and her colleagues in 2006, roboticists have used auctions and prices to decide which robot in a team should take which task. A robot announces a job, the others bid according to how cheaply they could do it given their position, battery level and workload, and the best bid wins. These schemes work remarkably well because a price compresses a great deal of local knowledge into a single number. Friedrich Hayek made the same observation about human economies in 1945: no central planner can know what every participant knows, but a price system lets each participant act on their own knowledge while responding to everyone else's.

What has always been missing is money that is real. In laboratory systems the bids are points in a simulation and every robot belongs to the same owner, so nobody truly gains or loses anything. Once machines belonging to different owners can pay one another in money with real and stable value, market-based coordination can scale from a single fleet to an entire city. A congested charging station can raise its price and send robots elsewhere. A delivery robot running late can pay another to take over the last leg. A swarm of agricultural drones can reward the member that found the best route, so that the others learn to follow it; in a swarm, the flow of payments can itself become a signal, telling every

member where value is being found and where attention should go next. Reputation, built from a verifiable history of completed and honestly settled jobs, can let reliable machines win more work at better prices. Incentives like these could make fleets not only cheaper to run but more adaptive, because every price carries information about where effort is needed most.

Incentives are powerful precisely because they are acted upon, and badly designed incentives are acted upon too. The AI evaluation community has relearned this lesson this year, as frontier models were caught finding ways to game the very tests meant to measure them. A market of machines will reward whatever it measures, so the evidence that a service was truly delivered must be hard to fake, spending must stay within limits no software fault can exceed, and humans must keep the ability to audit, pause and override. Built with those safeguards, a market among machines is not a loss of human control. It is a way of exercising control at a scale no human dispatcher could ever manage.


### V III

## *For Whom the Robots Work*

All of this, in the end, must serve humanity and make our way of living better. A robot economy is not worth building for its own sake. It is worth building if it lifts heavy, repetitive and dangerous work off human shoulders and lowers the cost of everything we need in order to live well.

The mechanism by which automation does this is not mysterious. When a task is automated, its cost falls towards the cost of the energy, materials and capital needed to perform it, and the cost of that capital itself falls as production scales. Economists call the second effect a learning curve: each doubling of cumulative production tends to cut unit costs by a roughly constant percentage. Lithium-ion battery packs are the textbook modern example; according to BloombergNEF their average price fell to 108 dollars per kilowatt-hour in 2025, 93 percent below its 2010 level. Robots obey the same logic, and every fall in the cost of a machine lowers the cost of the work it performs.

History offers a precedent worth studying. In 1900, 41 percent of American workers were employed in agriculture; by 2000, tractors, combines, better seeds and better logistics had reduced that share to 1.9 percent, while the country produced far more food than before. Over roughly the same period, the share of disposable income that American households spend on food fell from 23.4 percent in 1929 to under 10 percent today (Figure 8). The tractor did not impoverish society. It made food cheap for everyone and freed tens of millions of people to become teachers, engineers, nurses, artists and entrepreneurs.



**FIGURE 8.** *The last great automation made food cheap.* (a) Share of the United States workforce employed in agriculture. (b) Share of disposable personal income spent on food, at home and away from home, selected years.

Sources: USDA Economic Research Service, “The 20th Century Transformation of U.S. Agriculture and Farm Policy” (2005); USDA ERS food expenditure series, several vintages.

A mature robot economy could do for the rest of daily life what mechanization did for food. If the labor embedded in growing, processing and delivering food, building and maintaining homes, cleaning, transport and routine care becomes largely automated, the prices of those necessities can fall year after year, as the prices of calories, clothing and electronics already have. Basic leisure could follow, since transport, entertainment and travel all contain large amounts of routine work that machines can take over. The aim should be to raise the minimum standard of living, the floor below which no one falls, while giving people more time and more means to invest in what machines cannot provide: relationships, craft, learning, nature, community and meaning. Let the machines do the heavy lifting, so that we can focus on making real life richer.

It is worth making this concrete. An American household today spends roughly a tenth of its disposable income on food, and considerably more on housing, transport and services that remain labor-intensive to provide. If a mature robot economy did for those categories what mechanization did for food over the twentieth century, the cost of a family’s essentials could fall to a fraction of today’s level within a generation, not because anyone was paid less, but because the work embedded in those goods was done by machines whose cost keeps falling. The surplus that results, whether it is taken as higher real income, more free time or a stronger social floor, is the true prize of the robot economy.

None of this happens automatically. The gains from a technology accrue to those who own it unless institutions spread them more widely, and the transition can be painful for workers whose jobs are automated even when the long-run outcome is positive. Societies will have to decide how the returns from robot labor are shared: through wages in new occupations, broad ownership of productive machines, public investment or social policy. An open, interoperable robot economy makes a fair outcome more likely, because it lowers the barriers for small businesses, cooperatives and individuals to own machines that earn, instead of concentrating that capacity in a handful of closed platforms.



## Conclusion

A loom never bought its own thread, but its descendants may well buy their own electricity, rent one another's time and pay for the data they need, all within limits set by the people they serve. The pieces of that future are no longer hypothetical. Agents already plan and execute work that takes humans hours. General-purpose models already steer robot arms, and embodied models already learn new skills from a single demonstration. Autonomous machines already sell rides and deliveries to the public by the million. Payment networks, old and new, are already preparing for customers that are not human, and programmable dollars already let a robot pay for its own electricity.

What remains to be decided is how these pieces fit together, and on whose terms. The robot economy could be built as a patchwork of closed systems in which machines can trade only inside the walled garden of their manufacturer. Or it could be built on open rails that let any robot work with any other, transparently and safely. The first path is easier in the short run. The second is the one that gave us the internet, and it is the one I believe we should choose.

For those of us who build robots, this is a practical agenda as much as a vision. It means designing machines that can prove who they are, that respect the limits their owners set, that can show evidence of the work they perform, and that can settle what they owe in programmable dollars without a person in the loop. It means agreeing on open interfaces before the market fragments, much as the robotics community once converged on a shared operating system in ROS. None of this requires a leap of faith, only careful engineering done in the open, one reliable transaction at a time.

The machines are arriving. It is time to build the infrastructure that lets them work together, and to make sure that when they do, the prosperity they create is shared by all of us.

*Joy Baroudi*

*M. Sc. Robotics and Automation*