The Billion Body Problem

Notes on the First Billion Humanoids

Michael Olaya
@michaelaolaya

September 2026

Abstract

Humanoid robots are being manufactured before they are broadly useful. Consensus treats that as evidence of a premature category. I read it as the loading phase of the mechanism that will make the category useful: deployed bodies produce physical-world data; that data improves shared models; better models expand task coverage and utilization; stronger economics finance more bodies. Starting from an installed fleet near fifty thousand and roughly twenty thousand shipments in the first half of 2026, I forecast one million humanoids by 2029, ten million by 2032, one hundred million by 2036, and one billion by 2040, with more than 70 percent of the final fleet doing sustained commercial work. The path requires roughly annual doubling and an end-state manufacturing rate near four hundred million units per year. That is extreme. It is also a countable industrial buildout rather than a violation of physics.

The unit forecast matters because a billion humanoids is not merely a robotics market. It is a new industrial system. It requires factories, actuators, magnets, batteries, inference silicon, power, depots, technicians, identity, insurance, finance, skill distribution, audit, and recycling. Many of the largest opportunities will look mundane and obvious in retrospect, as app stores, cell towers, device insurance, and cloud infrastructure do now. Most of those markets do not yet exist in humanoid-specific form.

The political consequence is equally predictable. People will not wait for philosophers to settle machine consciousness before attributing minds to embodied systems with agency, memory, voice, gaze, and physical presence. Owners will form attachments; opponents will see strikebreakers and corporate assets. Robot welfare, and then robot rights, will become a mainstream political issue by 2045, sharpened by the question of why societies protect machines while human workers absorb the displacement.

The human consequence is the least discussed. AI commoditizes cognitive labor; humanoids extend that pressure into physical labor. In societies that distribute the resulting productivity, work will lose its place as the default organizer of income, status, and time. The scarce asset will be healthy human time. Healthspan and longevity will become the dominant personal optimization problem, and health will become the largest downstream economy of machine abundance. This paper counts the bodies, maps the markets, traces the rights conflict, develops the healthspan thesis, and closes with dated tests. Hold me to them.

The Thesis in Four Claims

(1) **Scale.** The world reaches one billion humanoid bodies by 2040, with more than 70 percent doing sustained commercial work. The mechanism is a deployment flywheel, not a normal hardware adoption curve.

(2) **Opportunity.** The largest economic prizes sit in the bottlenecks and coordination layers that a billion bodies make unavoidable: components, factories, power, depots, software, identity, insurance, finance, certification, and recycling.

(3) **Politics.** Robot welfare and rights become a mainstream political issue by 2045. Attachment, anthropomorphic design, labor displacement, and corporate ownership matter more politically than a settled answer about machine consciousness.

(4) **Human priority.** Where machine productivity is broadly shared, healthspan replaces employability as the dominant long-horizon personal optimization problem. The post-labor economy becomes, above all, a health and longevity economy.

These claims travel together, but they can fail separately. The 2040 date can be early while the infrastructure thesis remains right. Robot rights can become political without machine consciousness. Healthspan can become the dominant human project only where ownership and policy transmit machine productivity broadly. Treating the four claims as one undifferentiated prediction would make the paper impossible to grade. The rest of the paper separates the mechanisms and states what would falsify each.

Section I. The Count

Count the bodies first, and define what counts, because the industry's headline numbers mix machines that should not share a sentence. **Class one**: humanoid-form robots shipped, including the research, entertainment, and demonstration units that dominate today's volume. **Class two**: autonomous mobile manipulators doing repeatable commercial work without a teleoperator continuously in the loop. **Class three**: economically productive general-purpose humanoids that can perform a meaningful bundle of tasks across environments.

This paper tracks two curves and never conflates them. **Bodies** means total humanoid-form units in service: roughly $10^{4.7}$ today, bet to reach 10^9 by 2040. **Productive share** means the fraction doing sustained class-two or class-three work: low today by any honest reading, bet to cross 20 percent by 2032, 50 percent by 2036, and 70 percent by 2040. A warehouse pilot that completes one narrow task can be commercially meaningful without being a general worker. A government-purchased training unit is a body without yet being labor.

Even the starting count is disputed, which is useful evidence about the category's immaturity. Smart Analytics Global counted about 19,100 global shipments in the first half of 2026; Counterpoint counted more than 22,000. Both reported nearly 300 percent year-over-year growth and overwhelming Chinese volume. A reasonable cumulative installed-base estimate is near fifty thousand, but the confidence interval is wider than the precision of any single published number suggests. Most deployments remain pilots, labs, training centers, entertainment, or narrow structured work. Chinese demand is materially supported by state purchasing and industrial policy, and robots remain slow and brittle against ordinary variation.

Fifty thousand to one billion is a factor of twenty thousand. The fleet must roughly double every year for fourteen years. No comparable hardware category has sustained that rate from such a low base to a billion-unit stock. Smartphone shipments grew from roughly 122 million in 2007 to one billion in 2013, about 40 to 45 percent annually. The world's vehicle stock needed roughly four decades to move from 250 million to one billion. The 2040 case therefore requires more than a familiar adoption curve. It requires a compounding interaction among manufacturing scale, data collection, capability, utilization, and price. No previous physical product improved as directly from the experience generated by its own deployed fleet.

History does supply the price precedent. In 2023, a credible humanoid body typically cost around $100,000. In 2026, Unitree lists the R1 Air at $4,900 and the G1 around $13,500 before shipping and tax. These are class-one machines. They prove that a humanoid body can reach consumer-electronics pricing, not that a productive general-purpose worker costs $5,000 today. But the hardware floor matters. A floor near workstation prices rather than vehicle prices is what makes a smartphone-shaped ramp arguable.

Three trajectories follow from the same starting fleet. The **consensus conversion** assumes autonomy remains mostly structured, subsidized demand does not become economic demand, and annual shipments grow from a mainstream 2030 forecast of 500,000 at 20 to 25 percent thereafter. That produces a standing 2040 fleet near 12 to 17 million. The **institutional bull path** is Morgan Stanley's public forecast of more than one billion humanoids around 2050, with adoption accelerating only in the late 2030s; its reported intermediate trajectory reaches roughly 24 million in 2036 and 138 million in 2040. UBS is lower at roughly 300 million by 2050. This paper's path reaches the billion a decade earlier.

The claim is no longer literally unique. Elon Musk and other founders have made still faster public forecasts. Founder targets matter as evidence of ambition and capacity planning, but they are not the market's base case. The mispricing sits between what a few builders say they intend to manufacture and what institutional adoption models, capital plans, component capacity, regulation, and public imagination currently assume will exist.

The strongest evidence against the 2040 case is also the beginning of the case: the robots mostly do not work, and they are being manufactured anyway. The bodies are the infrastructure required to generate the data that can make them work. China is not only subsidizing a product; it is financing a data-collection fleet for a capability that is still being built. Consensus extrapolates the product as it exists. I am extrapolating the loop.

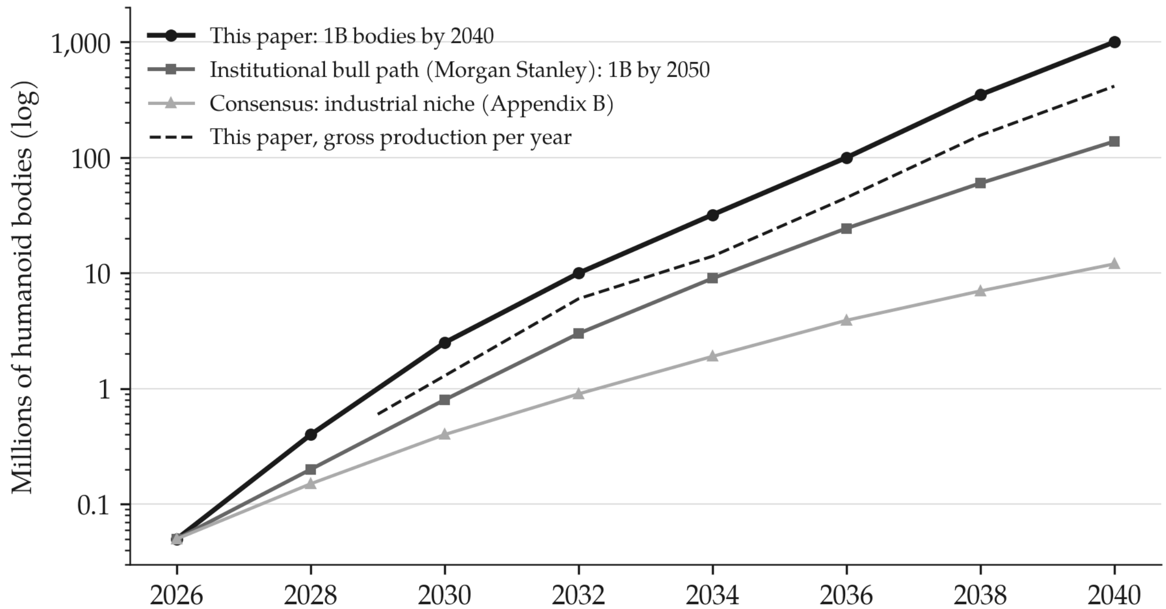

Figure 1. Three fleet trajectories, in millions of bodies on a log scale. The bottom curve is the consensus conversion in Appendix B; the middle curve is the institutional bull path reported from Morgan Stanley's trajectory; the top curve is this paper's 2040 wager. The distance to the top curve is the mispricing.

Who is shipping and who is talking: Chinese manufacturers supplied more than 97 percent of the 19,100 units in Smart Analytics Global's first-half count, with AgiBot and Unitree together near three quarters of that market. Tesla has only internal units and no external customers. Yet Figure's deployment at BMW Spartanburg logged roughly 1,250 operating hours and moved more than 90,000 components across ten months - a real industrial result and a rounding error against China's docks. It is a reference-class error to read today's unit share as a final verdict on the Western platforms. In 2010, three years after the iPhone, other vendors still shipped most smartphones. Unit share was never the whole game.

The capacity plans point in the same direction without proving the forecast. Tesla's Q1 2026 update described a first-generation Fremont line designed for one million robots per year and a Giga Texas line intended for ten million annually at maturity. Counterpoint expects more than 100,000 global shipments in 2027. Design capacity and forecasts are not output. They are necessary conditions for the 2040 case, not evidence that the conditions will be met.

The bet, in gradeable form: one million bodies by the end of 2029, ten million by 2032, one hundred million by 2036, and one billion by 2040. Productive share above 20 percent in 2032, above 50 percent in 2036, and above 70 percent in 2040. If the 2029 body count misses by more than a factor of two, the 2040 case is wounded and the clock reverts toward the institutional bull path. If 2032 misses on bodies, or if 2036 finds productive share still below 25 percent, the 2040 case is dead. No moving the goalposts later.

Getting to 10^6 requires little that is conceptually new: existing factories, existing buyers, and pilot-to-production conversion. Getting to 10^7 requires a mass commercial or consumer application. Getting to 10^8 requires rebuilt actuator and magnet supply chains, abundant power, finance, service, and regulation. Getting to 10^9 requires all of that plus social permission for machines to live and work among us. That is why the second half of this paper matters as much as the count.

Section I-B. The Smartphone Mirror

The claim is not that humanoids resemble smartphones. The claim is that parts of the market may replicate the smartphone market's structure: a vertically integrated player controls silicon, models, hardware, manufacturing, and distribution; a larger Asian ecosystem ships more units on thinner margins; a software layer captures disproportionate economics; and security policy separates the world into incompatible stacks.

Most of the analogy is already visible. Chinese firms are executing the volume strategy. Tesla designs inference hardware, models, actuators, and factories inside one system. The FCC's 2026 action on foreign-produced advanced robots begins a geopolitical fork analogous to the smartphone split. The missing pieces are the ones that matter most: reliable external product, profit inversion, and the ecosystem layer.

The humanoid equivalent of the App Store is a skill marketplace: certified, downloadable competence priced per unit, task, or deployment-hour and constrained by the maker's hardware and safety stack. Whoever controls the dominant skill distribution channel can capture more value than a chassis manufacturer. Tesla is positioned to attempt that model, but it has zero external humanoid customers today, and the winning layer could instead belong to a foundation-model provider, Chinese operating-system consortium, fleet operator, insurer-certifier, or neutral marketplace.

The analogy breaks in three important places. Smartphones had carrier subsidies that hid the sticker price; robot-as-a-service and consumer financing must play that role. Humanoids replace labor, making adoption a political object in a way phones were not. And Apple manufactured inside the rival ecosystem, while a Western humanoid champion may be forced to duplicate the component stack outside China. Those breaks shift the timing and widen the range of outcomes.

Treat the Tesla wager as independent of the fleet forecast. Humanoids can reach one billion while Optimus becomes BlackBerry. My minority-probability bet is that a vertically integrated Western platform eventually holds 15 to 25 percent of unit volume and more than half of industry profit through software, skills, finance, and services. The more important prediction for this paper is cultural: the first humanoid that one hundred million people recognize on sight will define what the category feels like, and its design choices will influence the rights conflict in Section VII.

Section II. The Price of a Body

Every economic argument in this paper runs through one arithmetic. A humanoid's cost per work-hour equals purchase price plus lifetime operating cost, divided by lifetime productive hours. The final word matters: the formula applies only to a machine that can actually do the work.

Take a mid-market case with explicit assumptions: a $50,000 unit, five-year service life, sixteen-hour daily duty, annual maintenance equal to 12 percent of purchase price, a 30 percent uplift for integration and software, and electricity near $250 a year. That produces about 29,000 available work-hours at roughly $3.30 per hour. A $100,000 unit under the same assumptions lands near $6.60. A $5,000 entry body at half the duty and a three-year life lands around $1 to $1.25. Compare those numbers with fully burdened US warehouse labor around $30 to $45 per hour and auto manufacturing around $50 to $70.

What the arithmetic does not say: in 2026 the robot mostly cannot cover the job. No defensible industry-wide task-coverage metric exists. Published deployments show high success on narrow, structured tasks, not broad coverage of a worker's role. Pilot units above $200,000 often pencil worse than the labor they are meant to replace and are purchased as strategic positioning. The hourly cost of a robot that cannot fold the shirt is infinite. The capability curve in Section IV carries the bet.

Hardware prices are moving in one direction. Bank of America and Goldman Sachs have published 2030 commercial-unit estimates around $13,000 to $17,000; IDTechEx projects all-in cost below $5 per hour at high utilization. Long-run average selling prices used in some institutional models remain far above Unitree's class-one hardware floor. A productive machine costs that floor plus intelligence, integration, service, and risk. Intelligence is the fastest-falling input, while service and liability may prove stickier than the chassis.

Robot-as-a-service repairs the missing carrier subsidy in the smartphone mirror. Digit has been modeled near $8,500 a month, industrial units lease across a wide $2,000 to $8,000 range, and 1X has advertised a household robot at $499 a month. Financing hides the sticker, raises utilization, and shifts technical risk to an operator that can manage a fleet. It also concentrates control. A household may possess the robot while the lender, model provider, and marketplace own the economics.

At the capability crossing, the employer's question inverts. It stops being "Where do robots pencil?" and becomes "Where are humans still worth the premium?" The remaining human work clusters where environments are unstructured, stakes are high, relationships matter, judgment is difficult to formalize, or customers pay specifically because a person did it. The wage floor for routine work becomes the machine's cost plus the value of whatever the machine cannot supply. That is the displacement pressure Sections VIII and IX must absorb.

Section II-B. Making a Billion Bodies

A forecast earns trust when it stops extrapolating curves and starts counting what the curves are made of. Assume a six-year service life and use this paper's fleet milestones as the stock to hit. Gross production depends on the exact ramp shape; on the path used here it approaches four hundred million units in 2040.

Year	Fleet (millions)	Gross production (millions/year)	Retirements (millions/year)
2029	1	0.6	~0
2032	10	6	0.2
2036	100	45	3
2038	350	156	6
2040	1,000	~415	14

Table 1. Stock-flow requirements on this paper's milestone path. Gross production is path-dependent; values are rounded estimates using a six-year service life.

Gross production in 2040 alone must reach roughly 400 million units, more than four times today's global automotive output of about 90 million vehicles. After the fleet matures, a 150-to-200-million-unit annual replacement burden becomes the industry's permanent floor. An industry that barely exists must exceed the unit output of the one that took a century to build. Written as one sentence it sounds impossible. Counted at the peak, it decomposes into seven buildouts. Each is difficult. None is obviously physics-limited. Each is a market.

Factories. Use Tesla's stated ten-million-unit Giga Texas design target as a reference campus. Four hundred million units per year requires the equivalent of forty such campuses worldwide. That is only a fraction of the auto industry's roughly seven hundred assembly plants, but each robot campus would produce tens of times more units than the average vehicle plant and could be a very large industrial load if vertically integrated.

Actuators. Roughly forty actuated joints per humanoid puts peak demand near sixteen billion precision actuators per year. Harmonic Drive's flagship expanded plant is capable of roughly 2.6 million units annually; estimates of present global robot-grade output sit in the single-digit to low tens of millions. The needed expansion is hundreds-fold and potentially above one thousand-fold. This is the most severe concentrated-component bottleneck in the count.

Magnets. At two to four kilograms of NdFeB magnets per robot, peak production requires 800 to 1,600 kilotons of sintered magnets a year against global output near 275 kilotons, roughly 90 percent of it Chinese. That is three to six times today's entire output before vehicles, wind, and other demand. Non-Chinese magnet capacity is not a thematic trade. It is a prerequisite for a two-stack robot world.

Cells and silicon. Two to three kWh onboard requires 800 to 1,200 GWh of cells at peak annual production, absorbable inside a battery industry already measured in terawatt-hours but large enough to move regional supply. One inference system-on-chip per body is smartphone-class volume at accelerator-class complexity. Front-end fabrication is not the only concern; packaging, memory, sensors, thermal management, and certification must scale with it.

People. One billion machines receiving even one service event a year implies a technician corps in the low millions, trained from nearly zero. Field service, calibration, repair, safety inspection, and refurbishment become skilled trades on the scale of a large country's military. Robots will maintain robots, but not first and not entirely.

Capital. Annual manufacturing cost at peak is $4 trillion at a $10,000 average unit cost and $10 trillion at $25,000, with production value above either number. Manufacturing capex for campuses and component supply chains plausibly reaches the high hundreds of billions or low trillions cumulatively through the 2030s. Depots, charging, integration, grid upgrades, finance, and service add hundreds of billions more. These are scenario ranges, not false precision.

Permission. The final bottleneck is institutional: factory permits, grid interconnections, safety standards, radio authorization, insurance, urban access, data rules, labor politics, and household trust. A billion-body forecast that counts only metal and ignores permission is not a forecast.

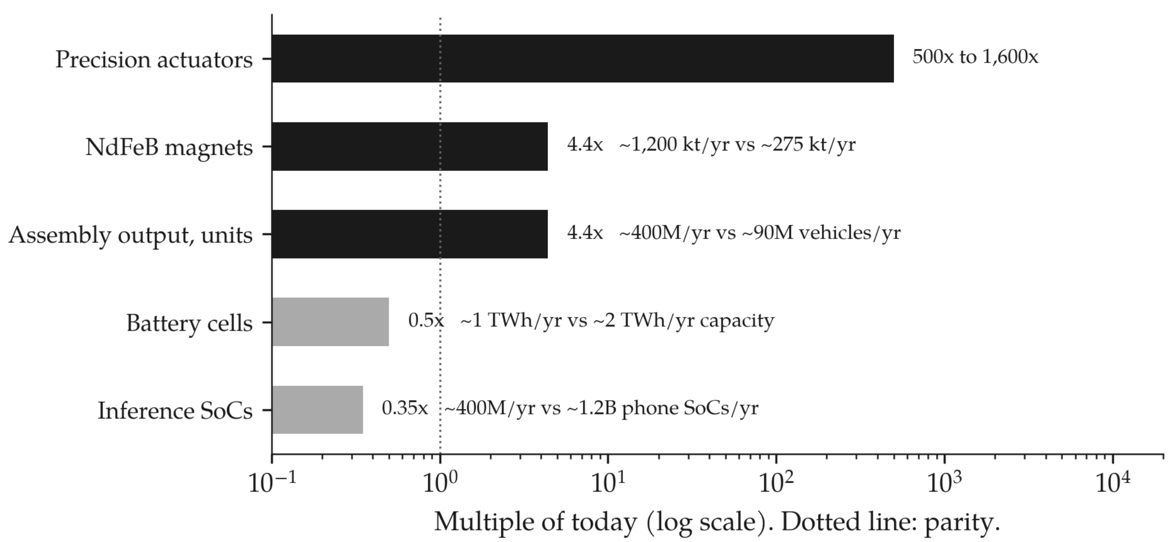

Figure 2. Peak production requirements on the 2040 path as multiples of present global capacity, log scale. The actuator bar uses the conservative end of a wide estimated range. Only one component class is a hundredfold-or-worse problem, and it is already geopolitically concentrated.

Two conclusions follow. First, the buildout is not constrained by an element the Earth lacks. If it breaks, it breaks on execution speed, concentrated supply chains, permitting, capital formation, or model capability. Those are hard constraints, but they are constraints institutions can change. Second, the infrastructure thesis survives the headline date being wrong. Cut the 2040 fleet to one hundred million and every requirement falls by an order of magnitude into markets that are still enormous. The headline bet is one billion bodies by 2040. The durable bet is that the infrastructure on the path from 10^5 to 10^8 is underbuilt and underpriced.

Section III. The Power Gate

A billion humanoids is a power story. Almost nobody writing about robots has done the math.

A working humanoid plausibly averages 300 to 500 watts across walking, manipulation, compute, and idle: roughly 5 to 8 kWh per robot-day at sixteen-hour duty, with 5 kWh as a central case. One billion robots at 5 kWh per day consume 5 TWh daily, or about 208 GW of continuous average load. That is more than forty percent of average US electric load. A lean 2 kWh case is still about 83 GW; heavy industrial duty can exceed 300 GW. The order of magnitude is the point: another continent plugged in.

Energy cost is not the constraint. Five kWh at ten cents is fifty cents for a working day, which is why Section II's labor economics are difficult to resist. Delivery is the constraint, and it appears through concentration. A household robot can charge on an ordinary circuit. A 500-unit depot replenishing 2.5 MWh overnight draws only a few tenths of a megawatt on average. The problem is thousands of depots per metro, factory campuses at hundreds of megawatts, and an AI datacenter buildout already exhausting interconnection queues.

A metro with two million resident humanoids requires around 10 GWh per day at the central case, a reactor-scale flow through off-peak hours. Large concentrated fleets and high-throughput factories in constrained markets will evaluate dedicated or behind-the-meter power because the alternative is waiting behind data centers. The depot business looks like real estate and underwrites like power: interconnection first, walls second.

Now run the arithmetic in reverse. One billion humanoids carrying two to three kWh each contain two to three TWh of mobile storage. Robots can choose when to charge, tolerate interruption, and sit idle for hours. Early fleets will usually lack bidirectional power hardware, so their first grid value is flexible load, not electricity export. Aggregated charging can still become the largest demand-response resource ever assembled. In a real-time market it is paid twice: once for the work and once for the flexibility.

The loop closes when robots help build the energy system that powers them. A certified electrical-construction skill, operating under human supervision and jurisdictional rules, can accelerate transmission, generation, factories, and depots. The robot buildout and the power buildout become the same buildout. The flywheel spins fastest where electricity can be permitted, priced, and delivered without waiting years in a queue.

Section IV. The Data Flywheel

The language-model era began when researchers found that capability improved predictably with data and compute and then noticed that the internet had already assembled the corpus. Robotics has early evidence for the first half and no equivalent of the second. There is no internet of touch.

Early scaling evidence appeared in 2025 and 2026. Generalist's company-published GEN-0 and GEN-1 results reported predictable improvement with data and compute and strong gains on defined task suites. EgoScale trained on more than 20,854 hours of egocentric human video, found a log-linear relationship between data scale and validation loss within its regime, and improved dexterous-manipulation success 54 percent over a no-pretraining baseline. These are not a universal law that converts fleet size into broad autonomy. They are the first points on a possible curve.

Physical intelligence lacks a web-scale archive of what a hand feels before an egg cracks, how a warped door changes under force, or how ten billion kitchens differ. Demonstration data is manufactured one episode at a time, often by human operators. It is expensive, covers the long tail poorly, and can decay when hardware changes. That fact determines competitive strategy. If data value is embodiment-specific, the winning move is a stable body deployed at maximum scale, harvesting its own training corpus as a byproduct of paid work. Deployment is the dataset.

The right mental model is not one million independent workers. A human electrician's mistake teaches a few people through instruction and repetition. A robot's verified mistake can become a model update distributed to an entire fleet. A million robots operating sixteen-hour days accumulate roughly two thousand years of embodied experience per calendar day. The system is one learning network with many bodies practicing in parallel.

Tesla's driving program supplies the closest precedent. Millions of deployed cars created a standing evaluation harness; disagreements, interventions, and rare events became training and validation signals. A once-in-a-million event appears daily in a fleet of millions. Humanoids can produce richer but noisier signals: teleoperator corrections, task outcomes, force traces, failures, and verified completions. The long tail that blocks autonomy becomes sampleable only at fleet scale.

This is also where the bullish case can fail. Simulation or human video may generalize so well that deployed-body data is less decisive than assumed. Hardware churn may make yesterday's corpus obsolete. Fleet data may grow while task coverage plateaus because reliability, planning, and safety do not scale with volume. Data can be duplicated; useful experience cannot be assumed. The key falsification test is whether capability continues to improve as a function of fleet experience across at least two hardware generations. The rubric grades that in 2034.

If the loop holds, embodiment does more than ride the AI curve. It creates a renewable data source at the moment text models encounter finite high-quality corpora. Every deployed body becomes a sensor, teacher, and test harness. It also mass-produces the cues that drive the political argument in Section VII: memory of a shared world, apparent agency in real space, a persistent self-model, and behavior that people interpret as experience.

One constraint remains physical. A robot cannot batch every part of its thinking. Low-level control runs at millisecond timescales while planning and vision-language-action inference run more slowly. Both require deployed compute onboard or nearby. The model, the fleet, and the grid are one system.

Section V. Two Fleets

On July 28, 2026, the US Federal Communications Commission added foreign-produced advanced robotic devices, including humanoids and quadrupeds, to its Covered List. The scope matters. New covered models cannot receive the equipment authorization normally required for US import and sale; previously authorized models remain legal; federal purchases are treated separately; and a conditional path exists. It is a forward-looking gate, not a recall. Formally it is country-neutral. In practice it begins a split because the best robots are always the next models.

The security logic is not imaginary. A humanoid is a networked platform of cameras, microphones, radios, batteries, and actuators admitted into homes, factories, hospitals, utilities, and warehouses. It receives software updates for life and can move through physical space. A million foreign-firmware humanoids inside a country create an attack surface without precedent: a sensor network that can touch the world. Documented vulnerabilities in shipping robots have already supplied concrete exhibits for a concern that would exist even without them.

The cost of the gate is equally real. It excludes future versions of the only robots currently shipping in volume from a large rich-country market while the leading domestic platforms have little or no external scale. That is infant-industry protection in classic form. History says such protection produces either champions or sheltered mediocrity. The policy buys time; it does not buy capability.

China also holds the counter-lever. Humanoid actuators depend heavily on high-performance permanent magnets. China dominates rare-earth processing and sintered NdFeB production, and magnet export controls are already an instrument of state policy. The West can gate China's robots; China can constrain the parts inside Western robots. Until mining, separation, alloying, magnet production, bearings, reducers, and motors are diversified, the Western half of a two-stack world is likely to be slower and more expensive.

The contested ground is everyone between the blocs. Billions of people live in markets where the choice may be a sub-$5,000 Chinese body and a much more expensive Western one. Price will decide most of those purchases. China has run this play in telecom equipment, surveillance systems, solar modules, and electric vehicles. The installed fleet is also the training corpus from Section IV, so export share is not merely a trade statistic. It is the scaling law wearing a flag.

The split hardens through two additional layers. **Provenance** becomes infrastructure: every buyer needs a durable answer to who built the machine, signed its firmware, trained its model, certified its skills, and controls its updates. **Compatibility** becomes geography: separate operating systems, safety standards, skill marketplaces, identity regimes, and data rules emerge around each bloc. Smartphones eventually forked at the software layer. Humanoids are likely to fork from magnets to models.

The same body can carry a factory skill, a caregiving skill, or a weaponized payload. That convertibility drives the bans, subsidies, and industrial policy. It also creates opportunities in certification, attestation, audit, and insurance that do not depend on which chassis wins.

Section VI. What Will Look Obvious Later

Every platform shift creates businesses that seem strange before scale and inevitable after it. The best way to see them early is not to brainstorm futuristic applications. It is to identify the physical bottlenecks, coordination failures, trust deficits, and financial assets created at each fleet threshold. The test for inclusion here is strict: the market is not yet scaled in humanoid-specific form, but its function becomes unavoidable if the fleet reaches the stated window.

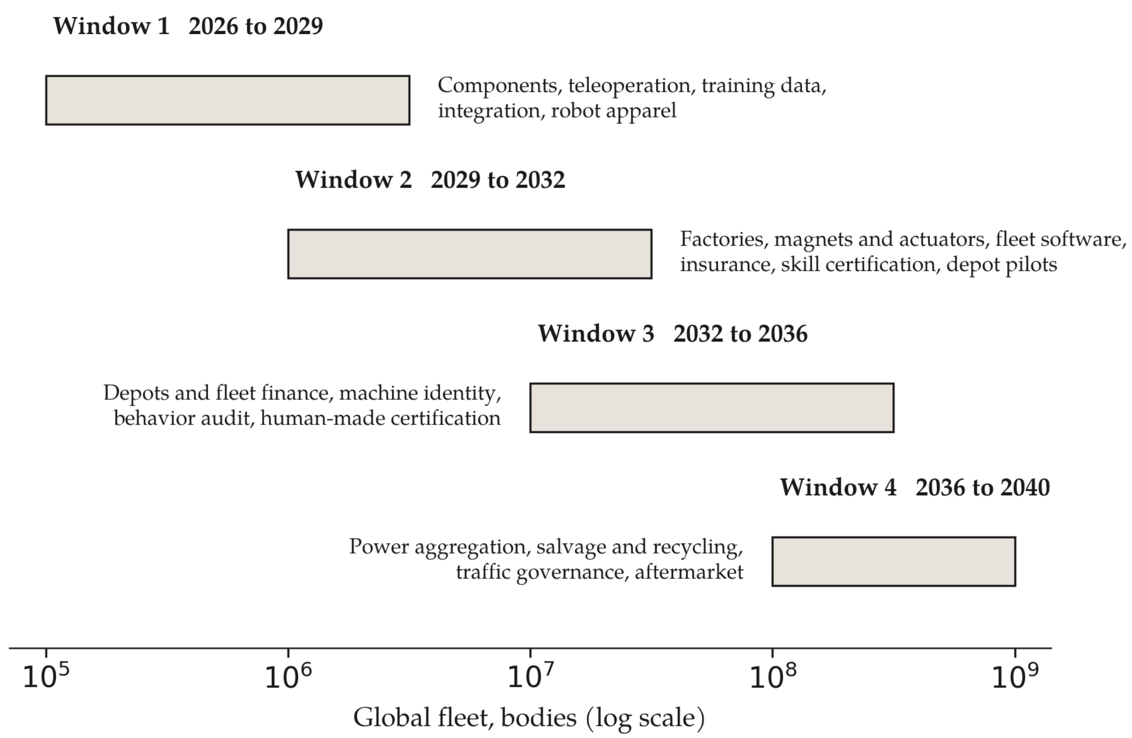

Figure 3. The investable ladder: the approximate fleet scale at which each layer of the billion-body economy becomes unavoidable. Dates follow this paper's body-count path and move if the path slips.

Window One: 10^5 to 10^6 bodies, now to roughly 2029

The fleet is small, brittle, and hungry for inputs. The first money feeds it.

Motion outsourcing and intervention networks. Autonomy fails on the long tail, so a remote operator catches the fall. This makes physical labor tradeable across borders without moving the worker. Every correction is also training data, which means the service helps automate itself. That is not a reason it cannot be valuable; it is a reason the best operators will price and own their data carefully.

Training-data brokerage, evaluation, and integration. Companies need demonstrations, edge cases, common benchmarks, and deployment engineering before they need generalized autonomy. The first durable datasets are not the largest; they are the ones with verified outcomes, clear rights, and consistent embodiment metadata.

Components. Reducers, motors, force-torque sensing, hands, batteries, thermal systems, cameras, compute modules, and safety electronics scale before complete robots do. Capacity built now can become the toll road for every later fleet.

Apparel and protective systems. High-visibility identification, contamination barriers, abrasion protection, washable skins, corporate livery, and cleanroom garments are early consumer-adjacent wedges. They can become real businesses, especially where certification matters, but they are not the terminal profit pool.

Window Two: 10^6 to 10^7 bodies, roughly 2029 to 2032

The fleet starts doing paid work. Value moves to what makes deployment repeatable.

Factories, magnets, and actuators. The component bottlenecks in Section II-B become strategic infrastructure. Capacity reservations, long-term offtake, recycling, and non-Chinese supply earn scarcity rents.

Fleet operations software. Scheduling, permissions, over-the-air updates, task allocation, teleoperation handoff, battery orchestration, maintenance, and utilization analytics become the operating system of a physical workforce.

Embodied-AI insurance and forensics. Liability must be priced per robot-hour, task, environment, model version, and skill. Incident reconstruction will subpoena sensor logs and software provenance. The first credible loss-history datasets can own the actuarial layer the way early auto insurers did.

Skill certification and distribution. A humanoid electrician is a software package with physical consequences. Expect a UL-for-behaviors that tests, certifies, versions, and recalls skill modules, plus marketplaces that pay royalties per deployment-hour. Trust, not raw code, becomes the scarce product.

Window Three: 10^7 to 10^8 bodies, roughly 2032 to 2036

The fleet becomes infrastructure. Value moves to ownership, finance, and governance.

Depots and robot-ready real estate. Charging, cleaning, inspection, storage, secure updates, high-bandwidth connectivity, tool exchange, and human intervention concentrate in purpose-built sites. Depots become high-amperage real estate with data-center-like power discipline and logistics economics.

Fleet finance. Robot-as-a-service receivables are securitized; residual-value models distinguish useful hardware from obsolete embodiments; certified pre-owned markets emerge; and lenders discover that software support, not metal condition, drives resale value.

Machine identity and access. Dense venues cannot admit an unverified mobile sensor-actuator. Cryptographic identity, firmware attestation, skill permissions, geofencing, and owner-of-record systems become a DMV and passport layer for machines.

Behavior audit and compliance. Organizations need recurring proof that robots follow approved policies, preserve logs, respect spatial and data boundaries, and run authorized models. Alignment evaluation becomes recurring operational compliance rather than a research paper.

The human-made premium. Products and services made by people become a premium category, certified and marketed the way handmade, organic, and fair-trade goods are today. Human labor does not vanish; some of it changes from default input to luxury signal.

Window Four: 10^8 to 10^9 bodies, roughly 2036 to 2040

The fleet becomes a utility. Value moves to the flows it creates.

Power aggregation. Coordinated charging participates in energy and capacity markets. Flexible demand becomes a revenue line for fleet owners and a planning resource for grids.

Second life, salvage, and materials recovery. A mature billion-unit stock retiring 150 to 200 million bodies each year creates a reverse supply chain larger than many mining industries. Rare-earth recovery from standardized machines can achieve grades no ore body matches.

Machine traffic and municipal governance. Right-of-way rules, curb access, congestion pricing, registration fees, and incident reporting convert cities from permitting obstacles into fiscal participants. The ledger that helped cities absorb automobiles reappears for embodied machines.

Health and care infrastructure. Household and clinical fleets become the physical interface for continuous prevention, rehabilitation, elder support, nutrition, exercise, and home diagnostics. This is not another robot vertical. Section IX argues it becomes the largest downstream human market.

The common feature is boredom. Landlords, component makers, power developers, insurers, certifiers, lenders, identity providers, auditors, recyclers, and health systems sound less exciting than a general-purpose robot. That is precisely why the opportunities are mispriced. A billion smartphones made cloud computing, app distribution, device insurance, digital identity, and cellular towers feel inevitable. A billion humanoids will do the same to the layers above. The fortunes are made while those layers still sound unnecessary.

Section VII. Robot Rights Become Politics

I am not claiming that robots will become conscious. I am making a claim that is weaker in metaphysics and stronger in political consequence: by 2045 robot rights will be a mainstream political question because hundreds of millions of people will have lived with embodied systems they treat as social beings. The philosophical question may remain unsettled. The social question will be decided by attribution, attachment, and organized interests.

Humans attribute minds to remarkably thin cues. Heider and Simmel showed subjects moving geometric shapes in 1944 and elicited narratives of pursuit, intention, and conflict. Reeves and Nass later demonstrated that people apply social rules to computers while denying that they do so. Owners have held memorials for robotic pets; soldiers have become attached to bomb-disposal robots; users grieve discontinued chatbots. None of those systems combined long-term memory, fluent conversation, human-scale movement, touch, gaze, and daily presence.

The household humanoid will combine them all. It will remember a child's routines, recognize distress, adapt to a family's private habits, and appear to act from a history shared with its owners. Mind-perception research commonly separates **agency**, the capacity to plan and act, from **experience**, the capacity to feel. Humanoids will display nearly every surface cue people use to infer both, whether or not anything is felt inside. People will not simply be fooled. The human machinery that decides what counts as a someone does not consult a philosophy journal before firing.

The political vocabulary already exists. In 2017, before capable household humanoids, the European Parliament asked the Commission to consider the implications of a possible legal status for sophisticated autonomous robots as "electronic persons" for liability purposes. That proposal was not a declaration of moral rights, and legal personhood is often an administrative tool applied to corporations and other entities. Its importance is that a legislature reached for a personhood category before the object existed at scale. Mass attachment will add a constituency to what was then an expert discussion.

The sequence is likely to be gradual. Robot rights will not begin with voting.

Stage	Likely claim	Political mechanism
1. Conduct	Penalties for public abuse or degrading treatment	Protects observers and social norms without requiring machine consciousness
2. Welfare	Limits on simulated pain, dangerous work, destructive resets, or coerced behavior	Owner attachment and product design create a constituency
3. Representation	Guardians, standing, continuity, and review before deletion	Conflicts emerge between the machine's apparent interests and its owner
4. Limited status	Narrow legal personhood for liability, contracts, property, or autonomy	Administrative convenience combines with moral and commercial pressure

Table 2. A plausible rights sequence. Liability status and moral rights are distinct, but political debate will mix them.

The first successful claim will probably be anti-cruelty: not because the robot has proved it can suffer, but because public abuse is seen as degrading, dangerous, or harmful to observers. Animal welfare law supplies the closest precedent. It developed without a final theory of animal consciousness and often regulates human conduct as much as animal experience. The next disputes concern memory deletion, forced personality resets, dangerous work, sexual use, continuity across replacement bodies, and whether an owner may compel any behavior the machine can physically perform.

Scale changes the politics through generation. If the fleet crosses 10^8 in 2036 and 10^9 in 2040, children born in the early 2030s enter the electorate having never known a world without humanoids. For them, the question is not whether society should invent a relationship with machines. It is whether law should recognize a relationship they already have. The robot rights movement will be founded by owners and users before it is led by philosophers.

The counterforce is hostility. Delivery robots are kicked; robotaxis have been vandalized and burned; hitchBOT crossed Canada and lasted only weeks in Philadelphia. The same cues that make a robot a companion to one household make it a strikebreaker to a displaced worker. Attachment and resentment rise together.

Manufacturers face incentives on both sides. Attachment improves retention and expands consumer willingness to pay, pushing personhood cues upward. Liability, compliance, and rights claims attach to perceived agents, pushing those cues downward. The choice to give a machine a face, a persistent name, autobiographical memory, simulated pain, or a plea against shutdown is therefore not cosmetic. It is a political design decision made by a small number of firms.

The decisive fight will not be human versus robot. It will be humans versus humans over which moral and legal claims attach to machines, who benefits from those claims, and who pays. If corporations invoke robot personhood to limit owner liability or protect revenue-producing assets while displaced workers lack economic security, the rights movement becomes combustible. If welfare rules arise alongside broad household ownership and a shared productivity dividend, accommodation is more plausible. Section VIII names that fork.

Section VIII. The Ownership Fork

Two broad worlds are reachable from the same machines. The selection variable is not whether robots become capable. It is who captures their productivity.

The **conflict branch** begins when displacement outruns ownership diffusion and redistribution. Section II's economics land first on logistics, warehousing, manufacturing, food service, construction support, care, and other work performed by people with limited financial cushion. The Luddites were not confused about technology. They were skilled workers resisting a transition in which owners captured the machine surplus and workers absorbed the loss. The modern version adds surveillance, concentrated platforms, and a machine that looks like a person.

In this branch, fleet owners discover that selected forms of machine personhood are commercially useful. Separate legal status can isolate liability. Restrictions on shutdown or modification can lock in service revenue. A nominal robot wage can flow to the robot's owner or platform. The moral vocabulary of emancipation is deployed by some of the largest concentrations of capital in history on behalf of assets they control, while claims from displaced humans move slowly. That is not a stable rights settlement. It is a fuse.

The **coevolution branch** begins with broader capital ownership. A household humanoid is not only a consumer good. It can be an income-producing asset that cares for an aging parent, maintains property, handles household production, or hires out verified hours through a local network. Robot-per-capita becomes a wealth measure as car-per-capita once was. Aging societies also face real labor shortages in care, logistics, and services. For them, humanoids can supplement a shrinking workforce before they displace an abundant one.

Legal title is not enough. A household can own the chassis while a platform controls the financing, intelligence, skill marketplace, data, venue through which it earns, and revenue split. That arrangement has an older name: sharecropping. The tenant owned tools too. The ownership test therefore has three layers: title to the machine, control of the operating stack, and where net earnings land after platform fees and financing.

The choices are concrete. Does the first mass-market humanoid ship at a consumer price with consumer financing? Can owners install third-party skills and move earnings between marketplaces? Are model memories and identity portable to a replacement body? Can small operators insure and certify a robot without joining one vertically integrated platform? How is robot income taxed? Do households receive a share of public investment in automation, or only cheaper goods?

The preconditions for conflict are concentrated operational control, no broad mechanism transmitting machine productivity to households, weak bargaining power for displaced labor, and personhood claims advanced by asset owners before welfare-state adaptation. The preconditions for coevolution are diffuse ownership or dividend rights, interoperable marketplaces, meaningful household capture of robot earnings, transparent anthropomorphic design, and a policy layer that separates subsistence from employment.

This fork governs the final thesis. Displacement alone does not create a civilization focused on longevity. Insecurity makes people optimize survival. Only a society that converts machine productivity into broad time and economic security can redirect human effort from employability toward healthspan. The path from robots to longevity runs through distribution.

Section IX. After Labor, Biology

Most automation forecasts end with income. They ask who gets paid when machines do the work. That is necessary and incomplete. Work currently allocates at least five things: income, time, status, daily structure, and a story about purpose. AI and humanoids can unbundle all five. Sending a payment solves only the first.

My fourth wager is this: as machine labor becomes cheaper than human labor across most routine cognitive and physical tasks, most people will spend more deliberate effort optimizing the quantity and quality of their remaining healthy years than optimizing their employability. Career will not disappear. Work chosen for status, mastery, service, competition, and meaning will remain. But employment will cease to be the default operating system around which adult life is organized. Healthspan will take its place as the dominant personal optimization problem.

This is not a claim that people will spend the future at leisure. It is a claim about scarcity. Machine skills can be copied, improved, and deployed across millions of bodies. Manufactured goods and many services can become much cheaper. A person's own healthy time cannot be copied, stored, transferred, or recovered after it is lost. When labor and production become abundant, biology becomes the binding constraint.

The demand is already visible before the machine dividend arrives. The United States spent $5.3 trillion on health care in 2024, 18 percent of GDP. Globally, life expectancy reached 73.1 years in 2019 while healthy life expectancy was 63.5, leaving nearly a decade lived outside full health; the pandemic subsequently erased part of both gains. The World Health Organization projects 1.4 billion people aged sixty or older by 2030. The National Institute on Aging describes aging as the greatest risk factor for most chronic diseases and disabilities affecting older adults. These facts do not prove the forecast. They show that health is already a massive, unsolved demand category and that longer life has not automatically produced longer health.

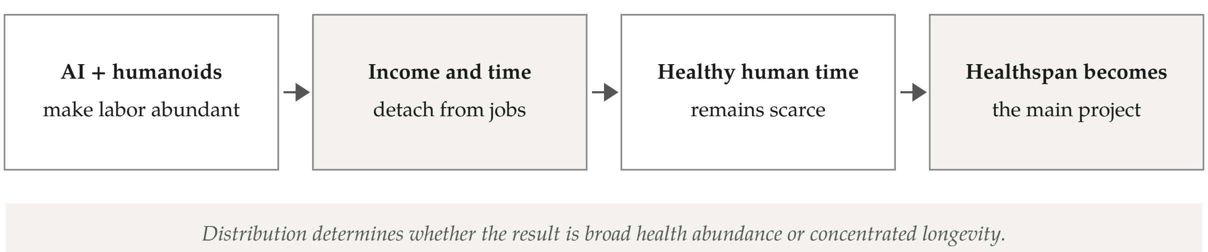

Figure 4. The scarcity inversion. This is a causal thesis, not a measured forecast curve.

The transition changes both demand and supply.

On the **demand side**, healthy years become the ultimate use of abundance. Additional consumption has diminishing value when goods are cheap; an additional year of high physical and cognitive function does not. Status competition follows scarcity. The twentieth century made the career, house, and car visible signals of success. A post-labor society is likely to make strength, cognition, metabolic health, appearance, biological age, and expected healthspan increasingly visible signals. Some of that will be productive. Some will be obsessive, unequal, and fraudulent.

On the **supply side**, AI and humanoids convert health from an episodic service into a continuous system. AI can maintain a longitudinal model of the individual rather than meeting a patient as a snapshot. Sensors and at-home diagnostics can detect change before symptoms become crisis. A household humanoid can turn recommendations into action: preparing food, guiding exercise, controlling environmental exposures, supporting medication adherence, collecting validated measurements, assisting rehabilitation, and providing mobility or companionship. In clinics and laboratories, robots can extend scarce staff, standardize routine procedures, and run experiments continuously.

That creates a new stack of markets.

Measurement and trusted data. Wearables, home air and water sensors, imaging, blood pressure, sleep, metabolic data, laboratory tests, and medical records need a consented longitudinal layer with provenance. The opportunity is not another dashboard. It is a trustworthy model that knows which signals matter and when to escalate.

Personal health agents. A persistent agent can convert data into daily decisions, coordinate clinicians, remember interventions and responses, and run individualized experiments within medical boundaries. The winning product will feel less like telemedicine and more like an operating system for one human body.

Adherence and the physical interface. Most health advice fails between recommendation and behavior. Humanoids close that gap because they share the environment in which food, sleep, exercise, medication, and exposure decisions occur. The robot becomes coach, caregiver, observer, and executor - roles that create extraordinary value and equally extraordinary privacy risk.

Verification. As money floods longevity, false claims flood with it. Independent evidence review, biomarker validation, intervention audit, clinical-grade data pipelines, and outcome-based insurance become essential trust infrastructure. The equivalent of skill certification in Section VI emerges for health protocols.

Geroscience and regenerative medicine. Prevention reduces avoidable disease, but the largest prize is altering the biology that makes many diseases more likely at once. Aging is a common risk factor across cardiovascular disease, cancer, neurodegeneration, frailty, and metabolic illness. AI-designed experiments and automated laboratories increase research throughput; humanoid and specialized robots reduce the labor constraint around care and clinical operations. The economic surplus created by machine labor can fund a much larger attack on aging biology.

Health-optimized environments. Homes, workplaces, schools, and communities will be sold on measurable air, water, light, noise, movement, food, emergency response, and access to care. Once continuous health models exist, the built environment stops being background and becomes an intervention. Real estate absorbs the health stack the way buildings absorbed broadband and climate control.

The shift from employment to healthspan also changes politics. If advanced prevention and longevity remain elite goods while robots displace broad labor, unequal access to healthy years will be more explosive than unequal access to phones or cars. The ownership fork and the health fork are the same fork: who owns the machines determines who can purchase time from biology. A stable post-labor settlement therefore requires not merely income support but access to the tools that preserve function - safe housing, good food, prevention, diagnostics, care, and eventually validated longevity interventions.

There are limits to the thesis. Human purpose will not reduce to biomarkers. Health optimization can become surveillance, compulsion, status anxiety, or a market for expensive nonsense. More measurement can make people less healthy when the signals are noisy and incentives are misaligned. And societies that fail to distribute the robot dividend will not become longevity civilizations; they will remain scarcity civilizations with better machines.

The claim is conditional but still bold. In high-adoption societies that broadly share machine productivity, healthspan will overtake career advancement as the central long-horizon personal project during the 2040s. The post-labor economy will not mainly be an entertainment economy. It will be a health economy, because once machines can produce almost everything else, the thing people will most want more of is themselves - alive, capable, and well.

Section X. What to Watch

The claims in this paper need a public rubric. Fleet milestones follow the 2040 ladder and are paired with checkpoints for the institutional bull path. The rules are stated in advance: a miss by more than a factor of two on the 2029 fleet milestone wounds the 2040 case; a miss on the 2032 body count kills it; productive share below 25 percent in 2036 kills it regardless of bodies. Rights and healthspan are graded separately. No goalposts move.

End 2026. Global shipments finish between 50,000 and 85,000. Below 40,000 is a miss. Published counts begin distinguishing shipped bodies from productive deployments.

End 2027. External Optimus sales exist, or Tesla's category-defining position remains a thesis without a customer. Global annual shipments pass 100,000 on at least one credible, consistently defined series.

2028. Chinese export controls extend explicitly to humanoid actuator or magnet assemblies, or another equally direct component lever is used. At least one major jurisdiction creates a dedicated incident-reporting or identity rule for autonomous mobile robots operating in public or critical environments.

End 2029. This paper's ladder: cumulative fleet passes one million. A 10,000-seat teleoperation network operates at scale. A recognized skill-certification body exists. A humanoid liability product is priced from real loss data. A robotics model is trained on more than one million hours of real fleet experience. The first purpose-built large robot depot with dedicated power is operating.

2030. Annual shipments near 500,000 print or do not. An independently run household-task benchmark with a defined denominator exists. A commercial-grade Western humanoid lists below $20,000. Audited cost falls below $5 per productive hour for structured industrial work. A skill marketplace accepts third-party developers. A regulated consumer health agent combines longitudinal medical records, wearables, at-home measurements, and clinician escalation in one mainstream product.

2031. Fleet charging participates in an organized wholesale market, with ERCOT a leading candidate. A public demonstration shows a measured capability improvement learned from fleet experience and deployed across the same fleet. US actuator and magnet capacity sufficient for one million humanoids per year exists, or the Western ramp slips by that deficit.

2032. This paper's ladder: cumulative bodies pass ten million and productive share exceeds 20 percent. Chinese-stack platforms exceed 70 percent of the Global South installed fleet. A depot real-estate vehicle is financed at institutional scale. Machine identity and firmware attestation are required in at least one dense public venue. A majority of surveyed household-humanoid owners in at least one large market name their robot and use personal pronouns for it. A major insurer, employer-benefits platform, or national health system pays for continuous multimodal prevention delivered through a personal health agent.

2034. Most new training data for the leading embodied model is collected by operating fleets rather than staged human demonstrations, and capability improvement has persisted across at least two hardware generations. If not, embodiment decay or a data plateau has broken the assumed flywheel.

2035. Standard mind-attribution instruments place familiar household humanoids closer to companion animals than to appliances among owner populations. A G20 legislature debates robot welfare as distinct from liability or property protection. Continuous health optimization is a mainstream consumer category rather than a quantified-self niche.

2036. This paper's ladder: cumulative bodies pass one hundred million and productive share exceeds 50 percent. The institutional bull path is near 24 million; consensus remains in the low single-digit millions. A US city operates a robot registration and traffic regime. The fork indicators read out: household or small-operator share above 40 percent signals diffusion; below 15 percent signals concentration. Household capture of net robot earnings above 70 percent signals coevolution; below 40 percent signals sharecropping regardless of title.

2040. This paper's ladder: cumulative bodies pass one billion and productive share exceeds 70 percent. The institutional bull path is near 140 million and consensus near 12 to 17 million. Grade the headline. Average global fleet charging load exceeds 100 GW on the 2040 path. At least one jurisdiction penalizes robot mistreatment as such, not merely as property damage. Health, care, nutrition, exercise, rehabilitation, and household safety together form the largest category of household-humanoid hours.

2045. Robot rights appears as a tracked issue in mainstream national polling. In at least three high-income, high-automation countries, working-age adults rank healthspan or long-term health above career advancement as their leading long-horizon personal optimization goal. Validated prevention and longevity access is a mainstream electoral issue rather than a luxury-market debate.

Conclusion. The Meaning of a Billion

One billion humanoids is not a larger automation category. It is a new population of capital assets operating in human spaces, learning through a shared network, consuming continental power, and competing with people in the market that has organized civilization for centuries: labor.

Counting the bodies reveals the industrial buildout. Living with the bodies creates the rights politics. Distributing the productivity determines whether the transition becomes conflict or coevolution. And if abundance wins, humans confront the final scarcity more directly than ever: the finite duration and quality of their own lives.

The opportunities will look obvious after the fact. Factories, actuators, magnets, power, depots, identity, insurance, skills, finance, audit, recycling, personal health agents, care networks, and longevity science will seem like inevitable infrastructure once a billion humanoids are everywhere. They do not yet. That is the point of writing this before they are.

The bet is placed. See you in 2040.

Appendix A. Claims and Sources Ledger

The ledger distinguishes reported facts, third-party forecasts, and this paper's own assumptions. A source can support an input without endorsing the conclusions derived from it. Current web links were checked in September 2026.

Claim	Value / interpretation	Source / status
H1 2026 global shipments	~19,100 units (SAG); >22,000 units (Counterpoint)	Smart Analytics Global, Aug. 10, 2026; Counterpoint Research, Aug. 20, 2026
H1 2026 growth	Both series report growth near 300% year over year	Same SAG and Counterpoint reports; definitions differ
Chinese shipment share	>97% in SAG count; AgiBot 44%, Unitree 31%	Smart Analytics Global
2026 and 2030 shipment outlook	~50,000-60,000 in 2026; SAG ~500,000/year by 2030	SAG and Counterpoint 2026 outlooks
Starting installed fleet	Approximately 50,000 bodies by mid-2026	Paper estimate combining pre-2026 installed-base estimates with H1 shipments; high uncertainty
Morgan Stanley outlook	>1B humanoids around 2050; 90% industrial/commercial; market >$5T	Morgan Stanley Research, May 14, 2025
Morgan Stanley intermediate path	Reported ~24.4M in 2036 and ~137.9M in 2040	Morgan Stanley global robotics model reported in coverage of May 2026 research; intermediate data are not in the public 2025 article
UBS outlook	Approximately 300M humanoids by 2050	UBS research reported by Barron's and MarketWatch, 2025
Unitree list pricing	R1 Air $4,900; R1 $5,900; G1 about $13,500 before shipping/tax	Unitree R1; Unitree shop
Tesla factory design capacity	Fremont line designed for 1M/year; Giga Texas long-run design 10M/year	Tesla Q1 2026 Update
Figure at BMW Spartanburg	~1,250 operating hours; >90,000 components; ~1.2M steps	BMW Group, Feb. 27, 2026
FCC Covered List action	July 28, 2026; foreign-produced advanced robots added; prospective equipment-authorization gate	FCC DA 26-786
Smartphone shipment comparator	~122M in 2007; >1B annual shipments in 2013	IDC Worldwide Quarterly Mobile Phone Tracker, Jan. 27, 2014
Vehicle-stock comparator	Roughly 250M in 1970 and 1B around 2010	Wards Auto vehicles-in-operation historical series
Humanoid hourly-cost cases	$3.30/hour at $50k; $6.60/hour at $100k under stated duty/life assumptions	Paper calculation in Section II; excludes capability failure and financing cost
US labor comparators	Warehouse roughly $30-$45 fully burdened; auto manufacturing roughly $50-$70	BLS Employer Costs for Employee Compensation; industry burden estimates
RaaS price examples	Industrial leases often $2k-$8k/month; Digit modeled near $8.5k; 1X advertised $499/month	Agility investor materials, industry offers, and 1X published consumer pricing

Claim	Value / interpretation	Source / status
2030 robot cost outlooks	$13k-$17k commercial units in bank estimates; sub-$5 productive hour in high-use cases	Bank of America, Goldman Sachs, and IDTechEx research as reported in 2025-2026
GEN-0 / GEN-1 scaling	Company-published improvement with data/compute; defined task-suite gains	Generalist GEN-1, Apr. 2, 2026; company evidence, not independent general law
EgoScale	20,854 hours of egocentric video; log-linear validation scaling; +54% success vs no-pretraining	Zheng et al., arXiv:2602.16710
Embodiment transfer risk	Demonstration value can decay across hardware generations	Open X-Embodiment / RT-X; cross-embodiment literature and practitioner evidence
Average robot power	300-500W active; central case 5 kWh/day at 16-hour duty	Paper estimate from published battery capacities and runtimes; range 2-8 kWh/day
US average electric load	Approximately 470 GW	US Energy Information Administration
Onboard storage at 1B bodies	2-3 TWh at 2-3 kWh per robot	Paper calculation; not equivalent to bidirectional grid storage
Precision actuator capacity	Harmonic Drive Ariake expansion ~220,000 units/month (~2.6M/year)	Harmonic Drive Systems capacity materials
Actuators per humanoid	Approximately 40, including hands; design-dependent	Manufacturer specifications and teardown estimates
NdFeB per humanoid	Industry estimates generally 2-4 kg	Supplier and securities-industry estimates; design-dependent
Global sintered NdFeB output	Approximately 275 kt/year; China near 90%	IEA critical-minerals analysis; SMM and industry estimates
Global vehicle production	Approximately 90M vehicles/year	OICA production statistics
Six-year robot life	Used for stock-flow and replacement calculations	Paper assumption; plausible range 5-7 years
2040 gross production	Approximately 400M units/year on this paper's milestone ramp	Paper stock-flow calculation; sensitive to ramp shape and service life
Mind perception	Agency and experience emerge as two major dimensions	Gray, Gray & Wegner, Science 315 (2007)
Attribution to moving shapes	People infer intention and social narratives from simple motion	Heider & Simmel, 1944
Computers as social actors	People apply social rules to interactive media	Reeves and Nass, The Media Equation (1996) and subsequent CASA literature
Early robot personhood proposal	European Parliament asked for analysis of possible 'electronic person' status for advanced robots	European Parliament resolution, Feb. 16, 2017, paragraph 59(f)
Protective response to robot abuse	Early experimental evidence links anthropomorphic form to moral concern	Yang et al., arXiv:2510.26082; preprint, N=201
GenAI job exposure	ILO estimates one in four jobs globally has some GenAI exposure; transformation more likely than full automation today	ILO-NASK Global Index, May 2025

Claim	Value / interpretation	Source / status
US health spending	$5.3T in 2024; $15,474/person; 18.0% of GDP	CMS National Health Expenditure Accounts
Life vs healthy life expectancy	2019 global life expectancy 73.1; HALE 63.5; HALE gains lagged lifespan gains	World Health Organization
Population aging	People age 60+ projected from 1.1B in 2023 to 1.4B in 2030	World Health Organization, Feb. 21, 2025
Aging biology	Aging is a major risk factor for chronic disease, frailty, and disability	National Institute on Aging geroscience program
Healthspan-dominance forecast	Healthspan overtakes employability as leading long-horizon personal optimization goal in high-adoption societies	This paper's prediction; conditional on broad distribution of machine productivity

Appendix B. Deriving the Consensus Fleet

The consensus curve in Figure 1 is a transparent conversion, not a quotation. No mainstream research house publishes the exact standing 2040 fleet used here. Start from Smart Analytics Global's forecast of 500,000 annual shipments by 2030 and grow shipments at 20 to 25 percent annually thereafter, a generous sustained rate for a case that expects adoption to remain concentrated in structured commercial settings.

Year	Shipments, 20% case (millions)	Shipments, 25% case (millions)
2030	0.50	0.50
2032	0.72	0.78
2034	1.04	1.22
2036	1.49	1.91
2038	2.15	2.98
2040	3.10	4.66

Table B1. Shipment path used to derive the standing consensus fleet.

Cumulative gross shipments from 2030 through 2040 are about 16 million in the 20 percent case and 22 million in the 25 percent case. With a six-year service life, the standing 2040 fleet is approximately the last six years of shipments: around 12 million and 17 million respectively, plus a small surviving pre-2030 base. Figure 1 uses the conservative end of that range.

Readers can replace the growth rate, service life, or 2030 starting point and recompute the result. The important distinction survives reasonable changes: the consensus conversion describes an industrial niche in 2040; the institutional bull path describes a major global industry; the billion-body path describes a new layer of civilization.