Foreword — what this document is and how to use it

v1 layer-organised deep dive: 38 IBKR-investable names across the humanoid + general-purpose robotics stack, from rare-earth magnets at the bottom to humanoid OEMs at the top. Integrates PhotonCap May 19 2026 mapping.

Investment Research

2026-05-21

Contents

Foreword — what this document is and how to use it

This document is the layer-organised analogue of the AI Universe v2 piece for the robotics build-out. It sits *above* the tier-organised `Robotics-MASTER.md` (May 14, 2026 — 38 Photoncap-template essays sorted into BUY / WATCH / SKIP buckets) and represents the same essays organised by **where each name sits in the physical robot**, not by where each name sits on the conviction ladder. The intent is to give the reader a way to walk the robotics value chain from the rare-earth mine at the bottom to the humanoid OEM at the top — which is also, not coincidentally, the order in which capital flows through the cycle.

PhotonCap, whose research methodology this workspace's Photoncap-style template was derived from, published "20 Companies in the Humanoid Cycle: Rare Earth Magnets to Vision-Language-Action" on May 19, 2026. That piece is the most rigorous external mapping of the humanoid supply chain published this year. It anchors on two events from the same week: Figure AI's Helix-02 livestream, in which three robots autonomously sorted 101,000+ packages over 81+ continuous hours with zero failures, and Schaeffler's announced 1,000-2,000 humanoid deployment programme with a five-year seven-digit-USD actuator supply contract with Humanoid AI (UK). Eighteen of his 20 names overlap our universe; the two non-overlaps (Inovance, Estun) are mainland-China A-shares deliberately excluded on the IBKR-tradability constraint. His mention-count signal — Schaeffler 46×, Tesla 20×, MP Materials 13×, NVIDIA 4× — is integrated throughout this document; the section that follows the stack overview is dedicated to walking through his GVM (Global Value Map) framework slot by slot.

Three pieces of context before the map.

One — the cycle is structurally early, but no longer hypothetical. As of late May 2026, two distinct commercialisation signals are now on the record. Figure AI's Helix-02 livestream demonstrated a humanoid program executing real production-line work for ~3.5 days continuously without intervention, a step-change in operational reliability. Schaeffler's Humanoid AI contract is the first publicly disclosed multi-year, single-

customer humanoid actuator supply deal at scale — the kind of off-take commitment that pulls supply-chain capex forward. Both events post-date the Robotics-MASTER write-ups by less than a week and validate the basket's central thesis: the suppliers monetise the build-out before the OEMs do.

Two — the supply-chain shape is monopolistic at the bottom, oligopolistic in the middle, fragmented at the top. Layer 4 (materials) has a single Western pure-play (MP Materials). Layer 1 (joint mechanics) has a Japanese oligopoly with Harmonic Drive at ~60-70% strain-wave share and Nabtesco at ~60% RV share. Layer 3 (compute) has a NVIDIA/Qualcomm duopoly at the cortex tier. Layer 2 (perception) is moderately fragmented but with named leaders per modality (Cognex/Keyence for 2D vision, Ouster/Aeva for lidar, VPG for tactile). Layer 5 (OEMs) is where competition is fiercest and concentration lowest — 10+ credible humanoid programs heading toward 20+. The implication for portfolio construction is symmetric: weight up where supply is concentrated, weight down where competition fragments.

Three — the visual style is intentional. The five-layer stack mirrors the eight-layer AI stack in `AI-Universe-v2-MASTER.md`. The framework is the same: walk down the cascade, identify the choke point at each layer, ask whether the valuation has absorbed the AI-mix shift. The names with the most concentrated supply *and* the lowest forward multiple are the trades. The names with the most fragmented supply *or* the richest forward multiple are the options. The map below makes the geometry visible.

Section 1 — The five-layer robotics stack

The robot is an electromechanical system that perceives, reasons and acts in a continuous loop. Every robot ever built — industrial arm, warehouse AMR, surgical robot, humanoid — runs that loop. The difference between a factory welder and a Tesla Optimus is not category but **degree**: how many actuated joints (a six-axis arm has 6; an Optimus has 28; an Atlas has 50+), how rich the perception, and how general the reasoning. Because the loop is universal, the suppliers in this universe sell into all of these end markets simultaneously. They are not betting on which humanoid OEM ships the most units; they are betting on the entire pie growing.

The stack has five layers, ordered from materials at the bottom to integration at the top:

Layer	What it covers	Key bottleneck names	Section
5 — OEM / system integrator	The companies that assemble all layers below into a finished robot. Worst structural margin position — buys from concentrated suppliers, integrates, competes on price	Tesla, UBTECH, Hyundai/Boston Dynamics, Rainbow/Samsung, Fanuc, Yaskawa, ABB, Symbotic, Teradyne	§7

Layer	What it covers	Key bottleneck names	Section
3 — Brain: compute + memory	Three sub-tiers: high-level reasoning at ~10 Hz (Jetson Thor, Dragonwing IQ); mid-level perception at 30-60 Hz (vision SoCs, FPGAs); motor control at 1-10 kHz (MCUs); plus state-save memory (MRAM)	NVIDIA, Qualcomm, Lattice, Ambarella, Ambiq, Everspin	§6
2 — Senses: vision, depth, touch	Three modalities — 2D machine vision (industrial leaders), 3D depth (li-dar/radar/stereo), and tactile (the dexterity bottleneck) — plus vision-processing SoCs	Cognex, Keyence, Mobileye, Ouster, Aeva, Arbe, Ambarella, indie, VPG, Allegro, Sensata	§5
1 — Body: joints, actuators, reducers	The mechanical cost spine. Each joint = motor + reducer + bearings + encoder + force sensor. Strain-wave (harmonic) gears used in wrists/elbows/fingers; RV cycloidal in hips/knees/shoulders	Harmonic Drive, Nabtesco, THK, Minebea, Nidec, Schaeffler, Allient, Novanta	§4

Layer	What it covers	Key bottleneck names	Section
4 — Fuel: materials	NdFeB rare-earth magnets, the irreplaceable input for every BLDC motor. China controls ~85% of NdFeB processing; no commercial substitute exists	MP Materials	§3

The order in the table is the order in which capital flows down the stack — the OEM at the top buys from the brain, the senses and the body, all of which depend on the fuel at the bottom. The order of the sections that follow is the reverse — we start at the bottom (the choke point) and walk up to the OEM (where competition is most fragmented).

Layer share of a humanoid bill of materials. A roughly $35-50k humanoid in 2026 distributes BOM cost approximately as follows, per Morgan Stanley Humanoid 100 (March 2026) and PhotonCap's GVM working numbers: joints (Layer 1) ~30%; sensing + perception (Layer 2) ~12%; compute + memory (Layer 3) ~10%; battery + power electronics (not in our universe) ~15%; structural frame, harnesses and assembly ~25%; magnets and other materials (Layer 4) ~8%. The implication is that Layer 1 is the largest single cost slice, and it is supplied by a tight Japanese oligopoly. That is where pricing power sits and that is why the joint layer earns the most attention in the per-name pages that follow.

Section 2 — PhotonCap May 19 2026 integration — the 20-company humanoid map

PhotonCap's "20 Companies in the Humanoid Cycle: Rare Earth Magnets to Vision-Language-Action" (May 19, 2026, paid) is the most concentrated third-party mapping of the listed humanoid supply chain published this year. The piece scores each name on the **Global Value Map (GVM)** — five axes: operating leverage, customer pull, technology moat, capacity headroom, and valuation gap. The mention-count signal in his text, which functions as a rough proxy for conviction-weighting within the piece, falls out as follows.

Tier A — anchor names (mention count ≥10). Schaeffler dominates the piece at 46×, anchored on the Humanoid AI five-year actuator supply contract; Tesla follows at 20× as the demand-side anchor via Optimus; MP Materials at 13× as the rare-earth magnet choke point.

Tier B — supporting names (mention count 4-9). NVIDIA at 4× as the compute-platform standard; Harmonic Drive Systems (6324.T), Nidec (6594.T) and UBTECH (9880.HK) at similar concentration in the joint-layer and OEM sections; Cognex (CGNX), Vishay Precision Group (VPG), Mobileye (MBLY), Ouster (OUST), Aeva (AEVA), Allient (ALNT), Novanta (NOVT) as the named sensing-and-motion suppliers.

Tier C — referenced names (mention count 1-3). ASMPT (522.HK), IQE (IQE.L) and a handful of Korean OEMs and ASIC-design names referenced in the cross-layer connections.

His twenty names overlap our 38-name universe at eighteen of the twenty (the two non-overlaps, Inovance and Estun, are mainland-China A-shares — both deliberately excluded from this basket on the IBKR-tradability constraint). The overlap is not a co-incidence — the PhotonCap framework and this basket's framework both derive from the same first-principles mapping of where pricing power sits in the value chain. What is interesting is the *weight distribution*. Our existing tier was 13 BUY / 18 WATCH / 7 SKIP, with Schaeffler at conviction 6 / Tier-2 WATCH. PhotonCap's 46× concentration on

Schaeffler reframes that allocation: the most disconfirmed-by-PhotonCap name in our top tier was the omission of Schaeffler from BUY. The integration call this document takes, and the only meaningful re-tier from the May 14 master, is **Schaeffler to Tier-1 BUY at conviction 7/10** on the strength of the Humanoid AI multi-year supply commitment. The 14× forward earnings multiple holds. The MP Materials, Tesla, NVIDIA and UBTECH calls in this basket are unchanged — his mention concentration on them corroborates rather than challenges our existing reads.

The cleaner read from his piece, beyond the Schaeffler upgrade, is the **convergence signal**: an independent third party with a distinct framework and a longer track record on the supply chain has converged on the same five names (MP, SHA, 6324, NVDA, TSLA) as the spine of the theme. Convergent calls from differently-derived frameworks are the highest-quality signal in this universe. They corroborate the basket's portfolio-rule recommendation: weight the bottlenecks (joint layer, magnets, premium sensing, NVIDIA), use OEMs as cheap optionality wrappers where valuation is subsidised by another business, reserve speculative sizing for pure-play humanoid OEMs.

The piece also embeds two new dated catalysts that did not exist on the calendar at the time the May 14 basket was written:

Event	Date	Direction	Touches
Figure AI Helix-02 livestream (3 robots, 101K+ packages, 81+ hours, zero failures)	13 May 2026 (concluded)	**Bull** — humanoid commercial-reliability proof	Indirect — every Layer 2/3 supplier with Figure designs (Mobileye is the visible one)
Schaeffler / Humanoid AI (UK) 1-2K humanoid + 5-year actuator supply contract	Announced ~16 May 2026	**Bull** — first multi-year humanoid actuator off-take	SHA direct; Harmonic Drive (6324) and Nabtesco (6268) read-through
BYD humanoid commitment — 20,000 units in 2026	FY26 ramp	**Bull** — Chinese demand ramp	9880 UBTECH (HK supply position), supply chain Tier-1 read-through
Tesla Optimus public-units target — 50,000-100,000 in 2026	FY26 ramp	**Bull / aspirational** — supply absorption	TSLA direct; magnet (MP) and joint (6324/6268) capacity gating

Per industry estimates published alongside PhotonCap's piece and Standard Bots' May 2026 industry mapping, global humanoid shipments hit roughly **13,000 units in 2025**, with Chinese firms — primarily Unitree and Agibot — accounting for ~80% of volume. The industry is on track to ship more humanoid robots in 2026 than in every prior year combined. The 2026 BOM cost is roughly half the 2024 BOM cost — driven by joint-layer cost reductions and compute commoditisation at the cortex tier. The implication is symmetric: cost is falling fast enough that the OEM-margin thesis is increasingly precarious (which is why our Tier-3 SKIP names look correctly de-rated), but unit volumes are growing fast enough that the supply-chain absolute revenue thesis is increasingly underwritten (which is why our Tier-1 BUY names look correctly Tier-1).

Section 3 — Layer 4 — Materials: the magnet bottleneck

Every brushless DC motor in a humanoid contains permanent magnets in its rotor. The magnet of choice is **neodymium-iron-boron (NdFeB)**, invented in 1984 and still the strongest commercially available permanent magnet by an order of magnitude. A typical Tesla Optimus contains roughly 28 BLDC motors plus a handful of fingertip micromotors, each with NdFeB magnets — on the order of 2-4 kilograms of NdFeB material per robot. China processes roughly 85% of the world's NdFeB. There is no commercial substitute: ferrite magnets are weaker by an order of magnitude; samarium-cobalt is more expensive and uses an even scarcer element. If an OEM cannot source NdFeB, it cannot ship robots.

Why this is the cleanest entry in the basket. As of May 14 2026, MP Materials is the only Tier-1 BUY name in this universe trading at Bucket A — within the entry zone, ~flat to its 50-DMA, RSI 44, undistorted by the broader robotics-momentum tape that has the rest of the universe printing RSI 70-90. PhotonCap's 13× concentration on the name corroborates this: it is the layer most aligned across independent frameworks. The fundamentals match: Apple's $72M magnet prepayment, GM's offtake agreement, the Department of Defense's funding commitment under the Defense Production Act — all three anchor a Western mine-to-magnet vertical integration that no other public name owns. The 2026 metal-and-magnet plant ramp is the visible catalyst.

The terminal risk for the layer is symmetric: a rare-earth-light or rare-earth-free motor architecture (Tesla has flagged exploratory work, but no commercial product exists), or a China-led NdFeB price war that destroys MP's pricing premium. Both are tail risks rather than central scenarios, and both would take years to play out.

MP — MP Materials Corp. · BUY (Tier-1) · Conv 7/10 · Bucket A

MP Materials Corp. (MP)

The only fully integrated US rare-earth-to-magnet producer — the NdFeB motor-magnet bottleneck behind every humanoid robot — backstopped by the Pentagon, Apple and GM, and the one name in this batch trading on a calm tape.

Investment Research · Photoncap-style deep dive · v1 of "MP Materials" · May 14, 2026

What MP Materials physically does

MP Materials sits at the chokepoint of the robotics supply chain that no compute or software company can engineer around: the motor magnet. Every humanoid robot, every robot arm, every actuator is built around electric motors, and the highest-performance electric motors use neodymium-iron-boron (NdFeB) permanent magnets — the strongest commercially available permanent magnets, which let a motor deliver more torque in a smaller, lighter package. A humanoid robot can contain dozens of motors; the magnet content is real, recurring, and physically irreplaceable with current motor architectures. The rare-earth elements at the heart of those magnets — neodymium and praseodymium (NdPr), plus heavy rare earths like dysprosium and terbium for high-temperature performance — are the binding material constraint, and China has dominated their mining, separation and magnet manufacture for two decades.

MP Materials is the Western answer to that dependency, and uniquely it spans the entire chain. It owns and operates Mountain Pass in California — the only operating rare-earth mine in the United States — where it mines and now also separates rare-earth ore into purified NdPr oxide and metal. It then runs the Independence facility in Fort Worth, Texas, where it produces rare-earth metal, alloy and — as of December 2025 — finished NdFeB permanent magnets. No other US company is fully integrated mine-to-magnet. MP describes itself, accurately, as the only fully integrated rare-earth producer in the United States covering the whole supply chain from mining and processing through metallization and magnet manufacturing.

Why is this the binding constraint of the robotics cycle in a way the chip names are

not? Because compute has multiple credible Western suppliers (NVIDIA, Qualcomm, Ambarella, Lattice) and the bottleneck there is competitive, not physical. Magnets are different: the material is geologically concentrated, the separation and magnet-making expertise is concentrated in China, and China has demonstrated — through its 2025 export controls — both the willingness and the ability to weaponize that concentration. If humanoid robots scale, the magnet supply chain is a hard physical constraint, and MP is the only listed pure-play that owns the entire Western alternative. The honest framing, developed below: MP's robotics revenue today is essentially zero — magnets only started in December 2025 and the customer relationships disclosed so far are Apple and GM, not robot OEMs — but the structural position is the most genuinely "bottleneck" of any name in this compute-and-materials batch.

Product roadmap

MP's roadmap is a vertical-integration build-out across two segments. The Materials segment is the established business: mining at Mountain Pass and the ramp of on-site separation, producing NdPr oxide and metal. Q1 2026 marked record NdPr output — the separation ramp is the near-term execution story. The Magnetics segment is the new leg: the Independence facility in Fort Worth commenced NdFeB permanent-magnet manufacturing in December 2025, and MP has said it expects initial magnet revenue in the second half of 2026. Beyond Independence, MP is building the "10X Facility" — a much larger magnet plant — and is advancing a magnet-recycling line at Mountain Pass tied to its Apple agreement, which would create an incremental domestic feedstock source containing both light and heavy rare earths. MP is also pursuing heavy-rare-earth expansion (dysprosium, terbium), the elements needed for high-temperature magnet performance and historically the hardest part of the chain to source outside China.

The dated milestones that matter: magnet manufacturing started December 2025; initial magnet revenue is guided for H2 2026; the recycling line has completed conceptual design and is moving through engineering and procurement; the 10X Facility is in build-out. What MP does not do is make motors, actuators or robots — it makes the magnet, the alloy and the separated oxide. It also does not yet have a disclosed robot-OEM customer; its anchor offtake relationships are in consumer electronics (Apple) and autos (GM), with the Department of Defense as a funding-and-offtake backstop. The robotics demand is a thesis about where magnet demand goes, not a current contracted book.

The financial print

MP Materials reported Q1 2026 on May 7, 2026: total revenue of $90.6 million, up 49% year-on-year, driven by higher NdPr oxide and metal sales as separation ramped and on stronger pricing. Including price-protection-agreement income, consolidated revenue plus PPA income was $132.9 million ($90.6 million revenue plus $42.3 million PPA income). Adjusted EBITDA swung to $36.6 million, a $39.3 million year-on-year improvement, helped by the revenue growth and the PPA income. The newly reporting Magnetics segment produced $21.1 million of revenue and $9.6 million of segment adjusted EBITDA — a real first contribution from the magnet business, though MP still posts a consolidated GAAP net loss as the build-out absorbs capital.

The funding structure is a core part of the financial story. MP received $32 million in Apple prepayments in February 2026, bringing total Apple-related prepayments on the balance sheet to $72 million. It has a Department of Defense partnership providing project funding and offtake support, GM prepayments, an $11.0 million Section 48C tax credit and 45X production credits — a stack of government and customer capital underwriting the Independence and 10X facilities and the heavy-rare-earth expansion. Sell-side analysts forecast roughly 67% full-year 2026 revenue growth with full-year EPS around $0.33; forward P/E is approximately 48.6x against a roughly $10.6 billion market cap.

Crucially, the technicals: MP prints an RSI of 44.3 and sits just +0.9% above its 50-day moving average. That is the single most important fact for the trade — MP is the only name in this entire compute-and-materials batch that is not extended. It has consolidated, not melted up, and that gives an entry an actual margin of safety the chip names and the memory names do not offer. The binary event is **Q2 2026 earnings, expected early August 2026** (MP reports on a roughly quarterly cadence; August 6, 2026 is the working estimate, to be confirmed), with the H2 2026 initial-magnet-revenue guidance the milestone to track.

Customer mix today

MP discloses by segment and by anchor relationship rather than by named-customer percentage. The Materials segment’s revenue is NdPr oxide and metal sales — historically a meaningful share of MP’s rare-earth concentrate went to Chinese processors, but MP halted all rare-earth exports to China as of April 2025, redirecting volume to its own domestic separation and to Western customers. The Magnetics segment, brand-new, has its offtake anchored by Apple (the $72 million of prepayments is tied to a magnet-supply agreement) and GM (which also has prepayments in place), with the Department

of Defense as the strategic funding-and-offtake backstop.

The 2024-to-2026 change is the entire story here. In 2024, MP was effectively a rare-earth concentrate miner whose output flowed substantially into the Chinese processing chain — a price-taking, single-product, geopolitically exposed position. By Q1 2026, MP has stopped exporting to China, is separating its own material domestically (record NdPr output), has started making finished magnets (December 2025), and has converted Apple, GM and the Pentagon into prepaying, capital-providing anchor partners. That is a transformation from "commodity miner exposed to China" to "integrated, government-backed, Western-supply-chain magnet producer." On robotics specifically: there is no disclosed robot-OEM customer. The robotics exposure is a demand thesis — robotics is forecast to become the single largest driver of NdFeB magnet consumption by 2040 per industry research — not a contracted relationship today. The honest framing: the customer mix is rapidly de-risking geopolitically and moving up the value chain, but the named customers are a phone maker and an automaker, with robotics as the structural tailwind behind the whole magnet build-out rather than a current line item.

What's actually happening at the Independence facility and the demand side

Two mechanisms to watch. On the supply side, the execution story is the Independence facility ramp: magnet manufacturing started December 2025, and the question is whether MP can scale finished-magnet output and convert the H2 2026 initial-magnet-revenue guidance into a real, growing line. The Q1 2026 Magnetics segment print — $21.1 million revenue, $9.6 million segment EBITDA — is the first proof point that the magnet business can contribute, supported by magnetic-precursor production while finished-magnet volume builds. The recycling line at Mountain Pass, tied to the Apple agreement, is the feedstock-security mechanism — completing engineering and procurement now, it would give MP a domestic source of both light and heavy rare earths independent of mined supply.

On the demand side, the mechanism is the structural shift in where NdFeB magnets get consumed. The industry research is consistent that robotics — and the humanoid sector specifically — becomes the dominant driver of magnet demand as production scales toward potentially billions of units over the long horizon, and global rare-earth magnet demand is already around 385,000 tonnes a year worth roughly $19 billion, growing in the high-single-digits. Be specific and skeptical, though: that demand is a long-horizon forecast (the "by 2040" framing matters), humanoid volumes today are tiny, and MP has not disclosed a single robot-OEM offtake agreement. There is also a

near-term demand nuance worth flagging — MP management itself, per Bloomberg's May 7, 2026 reporting, has signaled it sees heavy-rare-earth demand potentially falling as motor technology advances and magnet designs use less heavy rare earth. That is an honest complication: the company that should be talking its book on heavy-rare-earth scarcity is instead flagging that technology could reduce it. The robotics demand thesis for MP is real and large in the long run, but it is a structural tailwind, not a 2026 revenue catalyst, and the demand mix within rare earths is itself evolving.

The competitive threat / Lynas, USA Rare Earth and China

MP's competitive picture has a Western front and a Chinese front. On the Western front, the named competitors are Lynas Rare Earths — the largest rare-earth producer outside China, which in March 2026 began producing samarium oxide and is a commercial supplier of both light and heavy rare earths — and USA Rare Earth, which owns the Round Top deposit in Texas, is building a sintered-neo-magnet plant targeting first commercial production around Q1 2026, but likely will not open its mine until 2028. Against these, MP's edge is that it is the only fully integrated, currently operating, mine-to-magnet US producer — Lynas is stronger on heavy rare earths and is Australian; USA Rare Earth is earlier-stage and not yet integrated. There is no IP litigation defining the competitive timeline; this is a capacity-and-integration race.

The far more important competitive front is China. China dominates global rare-earth separation and magnet manufacture, has decades of cost advantage and scale, and demonstrated in 2025 that it will use export controls as policy leverage — which is precisely what created the Western-supply-chain investment thesis and the government backing for MP. But that same Chinese dominance is the structural threat: if Beijing chose to flood the market and crater rare-earth and magnet prices, MP — a higher-cost Western producer still ramping — would be squeezed, and that is exactly why the price-protection agreements, the government funding and the customer prepayments matter so much. They are the contractual and political backstop against a Chinese price war. The competitive reality is that MP wins on Western-supply-chain security and integration; it does not win on cost against China, and its economics depend on a structural willingness — by governments and by customers like Apple and GM — to pay a premium for non-Chinese supply.

The terminal risk

MP's terminal risk has two strands. The first is motor-architecture change: the entire magnet-demand thesis assumes high-performance electric motors keep using NdFeB permanent magnets. There are credible alternative motor designs — induction motors, switched-reluctance motors, externally-excited synchronous motors, and emerging rare-earth-reduced or rare-earth-free architectures — and if robot and EV motor designers move materially toward those to escape rare-earth dependency, NdFeB demand growth slows. The honest, uncomfortable data point here is that MP's own management has flagged that heavy-rare-earth demand could fall as technology advances — an acknowledgment from the most informed party that the demand mix is not fixed. The transition window is long (motor architectures change slowly) and NdFeB is unlikely to be displaced entirely because the performance-density advantage is real, but the magnitude of magnet demand per robot is not a constant.

The second strand is the China price war already described — Beijing using its cost-and-scale dominance to make Western production uneconomic. MP's defense against both is the same backstop stack: government funding and offtake, customer prepayments, and price-protection agreements that insulate the economics from spot-price collapse. The named beneficiaries if the terminal risks materialize are, on the motor-architecture side, the rare-earth-free motor technology developers, and on the price-war side, the Chinese incumbents. The terminal risk constrains the multiple — you cannot pay an unlimited price for a higher-cost producer whose demand thesis depends on both motor architecture and geopolitics holding a particular shape — but it is a slower-burning, more-hedged risk than the socket-absorption or category-absorption risks facing the chip and memory names, because MP's backstop stack is unusually deep and the strategic rationale for Western governments to keep MP viable is unusually strong.

Bull / Gap / Optionality (Photoncap framing)

1. MP is the only fully integrated, operating, mine-to-magnet US producer — a genuine physical bottleneck. Unlike the chip names, where the bottleneck is competitive, the magnet bottleneck is physical and geographic, and MP is the sole listed pure-play that owns the entire Western alternative chain — Mountain Pass mining and separation plus Independence magnet manufacturing (started December 2025).

2. The funding and offtake backstop is exceptionally deep. $72 million of Apple prepayments, GM prepayments, a DoD funding-and-offtake partnership, an $11.0 million Section 48C credit plus 45X production credits — a stack of government and blue-chip-customer capital underwriting the build-out and insulating the economics from a Chinese price war. Few small/mid-caps have a backstop this strong.

3. The financial inflection is real and visible. Q1 2026 revenue up 49% to $90.6 million, adjusted EBITDA swinging $39.3 million year-on-year to $36.6 million, record NdPr output, and a first $21.1 million Magnetics-segment contribution — the integration build-out is converting to numbers, with ~67% full-year 2026 revenue growth expected by the Street.

4. Uniquely in this batch, the tape is calm. RSI 44.3 and just +0.9% above the 50-day moving average — MP has consolidated rather than melted up, which means an entry here carries an actual margin of safety that the extended chip and memory names cannot offer. This is the cleanest entry setup in the compute-and-materials group.

5. The robotics demand thesis, while long-horizon, is structurally enormous. Industry research points to robotics — humanoids specifically — becoming the single largest driver of NdFeB magnet consumption over the long horizon, against an already-$19-billion, high-single-digit-growth magnet market. MP is the Western pure-play positioned for that demand.

Gap

1. Robotics revenue today is zero — it is a demand thesis, not a contracted book. Magnets only started in December 2025, initial magnet revenue is guided for H2 2026, and the disclosed customers are Apple and GM, not robot OEMs. The robotics tailwind is structural and long-dated; it is not a 2026 catalyst.

2. MP is a higher-cost producer that does not win on cost against China. Its economics depend on governments and customers continuing to pay a premium for non-Chinese supply. A Chinese decision to flood the market and crater prices would squeeze a still-ramping Western producer — the backstop stack mitigates this but does not eliminate it.

3. The demand mix within rare earths is itself evolving — per MP's own management. Bloomberg's May 7, 2026 reporting has MP signaling that heavy-rare-earth demand could fall as motor technology advances. When the company itself flags that the scarcity narrative may soften, the magnet-demand thesis deserves a haircut.

4. The valuation already prices a successful ramp. A ~48.6x forward P/E and a ~$10.6 billion cap on a company with a consolidated GAAP net loss requires the Independence and 10X build-outs to execute and the magnet revenue to scale on schedule. Execution slippage on the H2 2026 magnet-revenue guidance would hit the multiple.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~August 6, 2026 (to be confirmed)	Binary on separation ramp and magnet-revenue progress
Initial finished-magnet revenue	H2 2026 (company guidance)	Bull if delivered on schedule
First disclosed robot-OEM offtake agreement	Unscheduled	Bull — converts the demand thesis to a contracted book
10X Facility and recycling-line milestones	2026-2027	Bull — scales capacity and feedstock security
China rare-earth export-policy changes	Ongoing	Binary — tighter controls bullish, a price war bearish
Heavy-rare-earth expansion progress (Dy/Tb)	2026-2027	Bull — addresses the hardest part of the chain

---

The trade

MP is a Bucket A name for the robotics theme — and the standout entry setup in this compute-and-materials batch — because it is the only name here that combines a genuine physical bottleneck position, an exceptionally deep government-and-customer funding backstop, a real financial inflection, and a calm, non-extended tape. Initiate in a **$57-63 entry zone** (current $59.75 ± 5%; the calm RSI 44.3 / +0.9%-above-50MA tape means this is a true current-price entry, not a wishful pullback target), size at **1.5-2.5% of risk capital** — the margin of safety the tape provides supports a fuller position than the extended names in the batch warrant — with a **stop near $51** (below the 50-day moving average and the recent consolidation floor). The defining near-term binary is **Q2 2026 earnings around August 6, 2026**, with the H2 2026 initial-magnet-revenue guidance the milestone the whole magnet thesis hangs on. The honest caveat is that MP's robotics exposure is the longest-dated in this batch — zero revenue today, a structural demand thesis rather than a contracted book — so MP is owned for the bottleneck position and the inflection, with robotics as the long-horizon demand tailwind rather than the near-term driver. There is no cleaner listed pure-play for the Western rare-earth-magnet bottleneck — MP is the pure-play — so the only pivot is sizing: this is a core position in a robotics book precisely because, alone in this group, you are not

chasing it. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

Section 4 — Layer 1 — Body: the mechanical cost spine

The body layer — joints, actuators, reducers, bearings, encoders, force sensors — is the largest single cost slice in a humanoid BOM (~30%) and the most concentrated supplier base in the universe. A humanoid has 28-50 actuated joints. Each joint is the same five components in series: a brushless DC motor (fast and torque-poor) coupled to a reducer that trades RPM for torque, supported by 4-6 precision bearings, position-measured by an encoder, and force-measured by a strain-gauge sensor. The five components are made by five different supplier classes, all of them with high barriers to entry and an aging Japanese-dominated learning curve.

The two reducer technologies define the layer. Strain-wave (harmonic) gears, invented by Harmonic Drive Systems in the 1950s and still ~60-70% global share, use a flexible 200-tooth cup squeezed by an elliptical wave generator inside a rigid 202-tooth ring; each cam rotation advances the cup by exactly two teeth, giving 100:1 reduction in one stage with near-zero backlash. They are used in wrists, elbows and fingers — the precision-critical joints. RV (Rotary Vector) cycloidal gears, dominated by Nabtesco at ~60% global share, are a two-stage planetary-plus-cycloidal hybrid built for high load capacity; they are used in hips, knees and shoulders — the load-bearing joints. A humanoid uses both: roughly 24-30 strain-wave reducers and 6-10 RV reducers per robot. A 1,000-robot annual production run consumes 30,000-40,000 strain-wave reducers and 8,000-10,000 RV reducers — large numbers relative to Harmonic Drive's current humanoid-grade capacity.

The Japanese oligopoly is the structural feature of the layer. Harmonic Drive (6324.T), Nabtesco (6268.T), Minebea Mitsumi (6479.T), THK (6481.T) and Nidec (6594.T) sit at the top of five different sub-segments. Schaeffler (SHA0.DE) is the European cross-cut — bearings, compact gearboxes, integrated actuator units, with the most concrete humanoid customer commitment in the layer post-May 19. The two US public names — Allient (ALNT) and Novanta (NOVT) — are smaller and more diversified, with humanoid as an emerging optionality. Nidec is currently uninvestable on a live accounting scandal (third-party committee report February 2026; restatements back

to FY2022; FY25 results delayed; dividend forgone) and sits in the SKIP bucket — the structural opportunity is real, but the governance reset has to land first.

The catalyst stack inside the layer. The June 30, 2026 Schaeffler series-production start is the most concrete dated event. Harmonic Drive's FY-March 2026 results dropped in late April (revenue inflecting on robotics mix). Nabtesco's industrial-recovery thesis is reading through into 2026 prints. The PhotonCap May 19 piece pulled this layer into the spotlight by ranking Schaeffler at the top of its mention table — the implication is that the next tier of OEM-level commitments (BYD, Tesla, Figure expansion) will surface through this layer first.

6268 — Nabtesco / · BUY (Tier-1) · Conv 7/10 · Bucket B

Nabtesco (6268)

The RV cycloidal gear near-monopoly for heavy robot joints — a cleaned-up industrial compounder with a humanoid call option still ripening.

Investment Research · Photoncap-style deep dive · v1 of "Nabtesco" · May 14, 2026

What Nabtesco physically does

If Harmonic Drive Systems owns the compact wrist-and-finger joint, Nabtesco owns the heavy shoulder, base and elbow joint of the industrial robot — and it does so through a different gear architecture entirely: the RV (rotary vector) cycloidal reduction gear. Where a strain-wave gear flexes a thin steel cup, an RV gear uses a two-stage system: an input gear drives a set of crankshafts, which orbit a pair of cycloidal discs against a ring of pins. Because the cycloidal disc engages many pins simultaneously, the RV design carries far more torque for its size, shrugs off shock loads, and holds rigidity under the bending moments that a 6-axis industrial robot generates when it swings a 200kg payload. That shock tolerance and torque density is why essentially every medium-to-large industrial robot in the world has Nabtesco RV gears in its first three axes.

The scale of the position is the point. Nabtesco supplies precision reduction gears for roughly 60% of global medium-to-large industrial robot output, and the cumulative RV production base exceeds 14 million units. This is not a "robotics-adjacent" company — in the Component Solutions segment, Nabtesco is the binding-constraint supplier for the load-bearing joints of the global industrial robot fleet. The same cycloidal know-how also feeds machine tools, semiconductor manufacturing equipment, and positioners.

Nabtesco is broader than just robot gears, and the breadth has historically muddied the story. The group also makes aircraft flight-control actuators, railway brake systems, marine control equipment, automatic doors (it is a leading supplier of pedestrian automatic doors in Japan and Europe), and — until recently — hydraulic equipment for construction machinery. The 2025-2026 restructuring matters precisely because it sharpens this: management reclassified the hydraulic equipment business as a discontinued operation,

removing the most cyclical, lowest-return drag and leaving a cleaner industrial-precision and transport-infrastructure portfolio with the robot gear franchise at its core.

Product roadmap

The flagship is the Precision Reduction Gears RV series — the RV-N, RV-E and related families that span the torque range for industrial robot axes, with over 14 million cumulative units shipped. For collaborative robots and lighter applications Nabtesco built out the RVmini line, and for the highest-precision compact applications it sells the RD2 series and hollow-shaft variants.

The roadmap event directly relevant to this theme is the December 2, 2025 launch of two new compact-gear families: an expanded RVmini Series of compact, lightweight high-precision units, and the new Monocrank Series, which Nabtesco positions as delivering high precision and rigidity in a notably smaller package than a conventional RV gear. Both inherit the core RV cycloidal technology and are explicitly aimed at the new-application demand pool — collaborative robots and humanoid robots — where Nabtesco needs a physically smaller, lighter product than its industrial-robot heritage line to compete. Nabtesco also showcases digital actuators and exoskeleton systems, signalling an intent to move up the value stack from bare gear to integrated joint module. What Nabtesco does *not* make is the compact strain-wave reducer that dominates humanoid wrists and hands — that is structurally Harmonic Drive Systems' territory — so Nabtesco's humanoid entry is via miniaturised cycloidal designs competing for the higher-torque hip and knee joints.

The financial print

Nabtesco closed FY2025 (fiscal year ended December 31, 2025; results briefing dated February 19, 2026) with consolidated net sales of ¥307.9 billion, up 9.8% year-over-year, and operating income up 60.3% — a powerful operating-leverage swing in the right direction, the mirror image of what happened at Harmonic Drive Systems. Net income attributable to owners was ¥15,695 million, with basic EPS of ¥131.56. The March-quarter run-rate inside that print was strong: roughly 17% year-over-year sales growth, Component Solutions (the RV gear segment) up 28%, and operating margin up about 300 basis points to 9.9%. Crucially, this profitability improvement came *despite* reclassifying the hydraulic equipment business as a discontinued operation — so the clean-up did not cost earnings; it concentrated them.

Management is guiding to further earnings growth in FY2026 and has framed the year around shareholder returns and capital efficiency, consistent with the portfolio sharpening. At a JPY 5,704 share price and roughly JPY 668 billion market cap, the forward P/E of 52.9 is elevated for an industrial-gear maker — it embeds both the cyclical recovery already underway and a humanoid option premium. The binary is the Q1 FY2026 results, due around **May 8, 2026** — the print that tests whether the 28% Component Solutions growth rate is holding as the industrial-robot cycle and the early humanoid demand both build.

Customer mix today

Nabtesco does not disclose a clean named-customer table, so the mix reads by segment and end-market. The Component Solutions segment — the RV precision reduction gears — is the robotics core, selling into the FANUC, Yaskawa, ABB, KUKA tier of industrial-robot OEMs plus machine-tool and semiconductor-equipment builders; this is the segment that grew 28% in the March quarter. The Transport Solutions and aerospace-and-marine businesses (railway brakes, aircraft flight-control actuators, marine control) and the Accessibility/automation businesses (automatic doors, packaging automation) make up the balance and provide the stable, infrastructure-linked cash flow that funds the robotics investment.

Humanoid-specific revenue is still a small new-applications slice — Nabtesco itself frames humanoid and collaborative robots as an *emerging* demand driver for compact gears, not a current revenue pillar. Where Nabtesco can quantify a position, it is in the miniature RV reducer for humanoid robots category, where it is estimated to hold roughly 28% market share as of 2024 — a real foothold, but in a sub-segment that is still tiny in absolute yen. The structural shift to watch: in 2024 the investment case for Component Solutions was “cyclical industrial-robot recovery.” Through 2026, the December 2025 RVmini and Monocrank launches push the case toward “industrial recovery *plus* an early humanoid option” — but unlike Harmonic Drive Systems, the humanoid slice here is genuinely a call option, not yet a driver, and the honest framing is that the near-term print is an industrial-cycle print.

What’s actually happening at the industrial robot OEMs

The mechanism that drives Nabtesco’s near-term earnings is the industrial-robot capex cycle at its core customers, and that cycle turned up through 2025. The 28% March-

quarter growth in Component Solutions reflects industrial-robot OEMs rebuilding order books after the 2023-2024 automation-capex air pocket, with strength concentrated in the regions where Nabtesco's RV gears are spec'd into existing robot platforms. Because an industrial-robot model qualifies a specific RV gear and keeps it for the platform's life, Nabtesco's share at the established Japanese and European OEMs is sticky — the company is not winning new sockets so much as riding the volume of sockets it already owns.

The humanoid mechanism is different and slower. A humanoid OEM choosing a hip or knee joint reducer is making a fresh design decision, and Nabtesco's miniaturised cycloidal gears (the new Monocrank and RVmini) are competing against planetary gears, planetary roller screws and Chinese RV designs for those sockets. The ~28% estimated share in miniature humanoid RV reducers is encouraging, but it is share of a market that barely exists yet. The 2026-2027 humanoid design-win cycle is the period that will determine whether Nabtesco's compact-gear push converts the option into volume — and that is a 2027-2028 revenue event, not a 2026 one. Investors buying the 52.9x multiple need to be clear that they are paying today for a cash flow that is mostly still in the future.

The competitive threat / Zhongda Leader, Zhenkang and Shuanghuan

The competitive threat is the same structural story as at Harmonic Drive Systems, transposed to the RV gear. Chinese RV makers — Zhongda Leader, Zhenkang, and Zhejiang Shuanghuan — have crossed roughly 30% of Chinese domestic industrial-robot production by 2025, per industry data, and RV reducers from Zhongda Leader or Zhenkang run roughly 35-55% of Nabtesco's pricing for comparable torque ratings. Nabtesco is explicitly described in industry coverage as seeing its market share decline as these domestic players steadily build. The RV gear has historically been a harder process to replicate than the strain-wave gear — the cycloidal disc grinding and the crankshaft tolerances are demanding — which has given Nabtesco a longer-lasting moat than Harmonic Drive Systems enjoyed. But "longer-lasting" is not "permanent," and the Chinese cost gap is wide enough that, as with the harmonic reducer, Nabtesco's realistic future is keeping the precision and heavy-shock tier while ceding the price-sensitive tier.

The mitigant — and it is more substantive than at Harmonic Drive Systems — is that Nabtesco's RV franchise is anchored in the *installed industrial-robot base*, where requalification is expensive and Western and Japanese OEMs are conservative. The Chinese share gains are concentrated in Chinese-assembled robots; Nabtesco's grip on

the global, non-China OEM platforms is stickier. The bear case is that humanoid is a greenfield market with no installed-base inertia, so if Chinese RV makers compete for humanoid hip/knee sockets from day one, Nabtesco's option is contested before it is even exercised.

The terminal risk

The terminal risk for Nabtesco is twofold. First, the same margin-not-volume risk as Harmonic Drive Systems: Chinese RV capacity sets a global clearing price 35-55% below the Japanese level, and even if humanoid volumes arrive, Nabtesco earns a commodity margin rather than the monopoly margin its valuation implies. Second — and more specific to RV — is architectural displacement *within* the humanoid joint: the 2026 industry literature increasingly favours planetary roller screws and integrated linear actuators for the high-force hip and knee joints that would otherwise be RV-gear sockets, because roller screws offer a better power-to-weight ratio for linear-dominant leg motion. If humanoid legs go linear rather than rotary-cycloidal, Nabtesco's natural humanoid entry point narrows to a thin band of rotary leg and torso joints, and the company is left as a strong *industrial*-robot compounder with a disappointingly small humanoid attach. Nabtesco's credible defence is the Monocrank/RVmini miniaturisation push and the digital-actuator and exoskeleton work — but, as with Harmonic Drive Systems, the defence is unproven until the humanoid design wins actually print.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. A genuine ~60% monopoly in the heavy robot joint, now in a cleaner corporate wrapper. Nabtesco supplies RV gears for roughly 60% of global medium-to-large industrial robot output, with 14 million cumulative units shipped. The 2025-2026 hydraulics divestiture removed the most cyclical drag without costing earnings — the FY2025 profit improvement came *despite* the reclassification.

2. Operating leverage is swinging hard the right way. FY2025 net sales rose 9.8% to ¥307.9 billion and operating income jumped 60.3%, with the March-quarter Component Solutions segment up 28% and operating margin up ~300bps to 9.9%. This is the mirror image of the Harmonic Drive Systems trough — Nabtesco is on the recovery side of the same cycle.

3. **The industrial-robot capex cycle has turned and Nabtesco's sockets are sticky.** The 28% Component Solutions growth reflects OEMs rebuilding order books after the 2023-2024 air pocket, into RV sockets Nabtesco already owns on existing platforms. This is volume on installed share, not speculative new wins — the lowest-risk part of the thesis.

4. **The December 2025 RVmini and Monocrank launches are a real humanoid entry product.** Both inherit the proven RV cycloidal technology in a smaller, lighter package explicitly aimed at collaborative and humanoid robots. Nabtesco already holds an estimated ~28% share in miniature humanoid RV reducers — a foothold, not a standing start.

5. **The transport and accessibility businesses fund the robotics investment.** Railway brakes, aircraft actuators and automatic doors throw off stable, infrastructure-linked cash flow, letting Nabtesco invest in compact-gear capacity through the cycle without balance-sheet stress.

Gap

1. **Chinese RV makers are crossing 30% of Chinese robot output at 35-55% of Nabtesco's price.** Zhongda Leader, Zhenkang and Shuanghuan are taking domestic share, and Nabtesco is explicitly described as losing share to them. The RV moat is more durable than the harmonic-gear moat, but the cost gap is wide enough to cap the achievable margin.

2. **The humanoid slice is a call option, not a driver — and the multiple already pays for it.** Nabtesco itself frames humanoid as *emerging* demand. The ~28% miniature-RV share is share of a near-zero market. At 52.9x forward earnings, the price embeds humanoid revenue that is realistically a 2027-2028 event.

3. **Humanoid legs may go linear, not cycloidal.** The 2026 shift toward planetary roller screws and linear actuators for high-force hip/knee joints would narrow Nabtesco's natural humanoid entry to a thin band of rotary joints — routing the leg BOM to THK and screw makers instead.

4. **The near-term print is an industrial-cycle print, full stop.** If global industrial-robot capex stalls — a China slowdown, a Western capex pause, yen strength compressing export margins — the recovery leg that justifies most of the 52.9x multiple weakens, and the humanoid option is too small and too far out to cushion it.

Optionality

Event	Date / window	Direction
Q1 FY2026 results	~May 8, 2026	Binary on Component Solutions growth holding
Monocrank / RVmini humanoid design wins	2026-2027	Bull if a named humanoid OEM designs them in
Industrial-robot capex cycle trajectory	Through 2026	Bull/Bear — drives the near-term earnings
Hydraulics divestiture completion / capital return	2026	Bull if proceeds fund buybacks / robotics capex
Chinese RV share data (Zhongda Leader, Shuanghuan)	Ongoing	Bear if share crosses into precision tier

---

The trade

Nabtesco is the lower-beta, cleaner-balance-sheet way to own the joint-layer monopoly thesis — a real industrial-robot recovery you are paid for today, with a humanoid option you are paying a premium for but not yet receiving. That mix earns it Bucket B alongside Harmonic Drive Systems, but for opposite reasons: Harmonic Drive Systems is trough-earnings with the humanoid story already live; Nabtesco is recovering-earnings with the humanoid story still ripening. Initiate at JPY 5,419-5,989 (current JPY 5,704 ±5%; RSI 71.6 is momentum in an early-cycle theme, and the +22.2% vs 50MA is far less stretched than THK or Minebea), size at 1.5% of risk capital, and stop at JPY 4,750 (below the 50-day structure and the pre-breakout consolidation). The defining binary is the Q1 FY2026 print around May 8, 2026 — the test of whether the 28% Component Solutions growth rate is holding. If you want the same RV-monopoly thesis with the industrial-cycle leg already de-risked and a richer near-term earnings base, Nabtesco *is* the cleaner expression versus Harmonic Drive Systems; the pivot here is the other direction — if you want maximum humanoid torque and can stomach the trough multiple, Harmonic Drive Systems (6324) is the more levered call. **Conviction: 7 / 10.**

Sources referenced inline throughout: Nabtesco FY2025 results briefing material (February 19, 2026); Nabtesco RVmini / Monocrank product launch announcement (December 2, 2025); industry data on Zhongda Leader / Zhenkang / Shuanghuan Chinese RV share; market-research estimates on miniature RV reducer share for humanoid

robots. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

6324 — Harmonic Drive Systems / · BUY (Tier-1) · Conv 7/10 · Bucket B

Harmonic Drive Systems (6324)

The strain-wave gear near-monopoly — the core humanoid joint reducer, bought through a cyclical earnings trough as the secular volume story arrives.

Investment Research · Photoncap-style deep dive · v1 of "Harmonic Drive Systems" · May 14, 2026

What Harmonic Drive Systems physically does

Every articulated robot joint needs to do one thing: take a small, fast, low-torque electric motor and convert it into slow, powerful, precisely-controlled motion. The component that does this is a reducer (a gearbox), and for compact joints — robot wrists, cobot arms, and now humanoid shoulders, elbows and hands — the dominant architecture is the strain-wave gear, also called a harmonic drive. Harmonic Drive Systems did not just commercialise this device; the company's name *is* the generic term for it.

A strain-wave gear has three parts and no conventional gear train. A stiff outer ring (the circular spline) has internal teeth. Inside it sits a thin, flexible steel cup (the flexspline) with slightly fewer external teeth. And inside *that* sits an elliptical bearing-shaped plug (the wave generator) mounted on the motor shaft. As the wave generator spins, it flexes the cup into an ellipse, so the cup's teeth mesh with the ring's teeth only at the two ends of the long axis. Because the cup has, say, two fewer teeth than the ring, one full rotation of the motor advances the output by just two teeth — a reduction ratio of 50:1 or 100:1 in a single, flat, lightweight stage with effectively zero backlash. That combination — high ratio, no backlash, low weight, coaxial input and output — is exactly what a humanoid robot's distal joints require, and almost nothing else delivers it in the same package.

The "binding constraint" framing matters here. A humanoid robot built on the Tesla Optimus reference architecture uses roughly 28 actuators, of which around 14 are rotary; the rotary actuators in the arms, wrists and hands are overwhelmingly strain-wave designs. Harmonic Drive Systems sells the gear itself, the gear-plus-bearing unit, and increasingly the complete actuator (gear, frameless motor, encoder and driver integrated).

It does *not* make the RV cycloidal gears used in heavy industrial-robot base joints — that is Nabtesco's territory — and it does not make linear motion. It is a pure-play on the rotary precision joint, which is both its strength and its concentration risk.

Product roadmap

The core catalogue is the HarmonicDrive component series — the CSF/CSG/SHF families for general robotics and the CPL/CPU "pancake" units for flat joints — plus the CobotPro line launched in 2022 specifically for collaborative-robot integrators who want a drop-in gear-and-bearing module. Above the bare gear, Harmonic Drive Systems sells the FHA and RSF rotary actuators, which integrate an AC servo motor, the strain-wave gear, an encoder and a brake into one housing — the product the company wants humanoid OEMs to buy, because it captures far more bill-of-materials value than the gear alone.

The roadmap event that matters for this theme is the high-torque micro actuator co-developed with MinebeaMitsumi, unveiled for first public demonstration at CES 2026 (January 2026). It is a geared micro actuator measuring 13mm wide, 19.4mm high and 60.4mm long — sized for humanoid robot *hands and fingers*, the highest-actuator-count, highest-unit-volume part of the humanoid body. This is the clearest signal that Harmonic Drive Systems is positioning for the fingertip-level volume explosion rather than just the shoulder-joint count. Capacity is following: FY2026 capital investment is guided at ¥6,100 million, up 62% year-over-year, explicitly to build humanoid-driven reducer capacity. The medium-term plan, confirmed at the May 2025 results briefing, targets ¥90 billion in sales and ¥15 billion in operating profit by FY03/27 — a roughly 60% revenue increase over the FY2025 trough.

The financial print

Harmonic Drive Systems closed FY2025 (fiscal year ended March 31, 2025; results released May 20, 2025) with net sales of ¥55.65 billion, down 0.3% year-over-year, and — this is the headline — operating profit of ¥1,967 million, down 94.4% from ¥28,577 million the prior year. Operating margin collapsed from roughly 51% to 3.5%. Net income was ¥3.47 billion. This was not a demand failure; it was a savage operating-leverage unwind. The semiconductor-equipment correction hit the high-margin FPD/semi reducer business, the company kept SG&A and headcount up for the coming humanoid ramp, and the result was a near-total profit wipeout on roughly flat revenue. That is the trough the stock is being bought through.

The recovery is already visible in the run-rate. Cumulative 9-month FY2026 (April-December 2025) net sales reached ¥42.18 billion, up 4.5% year-over-year, with management citing surging demand from robot manufacturers in Japan, China and Europe plus a semiconductor-equipment recovery — and noting that orders to robot makers *exceeded* the earnings-forecast assumption. First-half FY2026 net sales were ¥27.84 billion, up 4.8%. The forward consensus, against the FY03/27 medium-term target of ¥90 billion revenue and ¥15 billion operating profit, implies operating margin rebuilding toward the high teens. At a JPY 6,430 share price and roughly JPY 608 billion market cap, the forward P/E of 85.0 is a trough-earnings multiple, not a steady-state one — the bull case is that FY03/27 normalised EPS makes the multiple look very different. The binary is the FY2026 full-year results, due on or around **May 19, 2026** — the print that confirms whether the robot-order momentum is translating to margin.

Customer mix today

Harmonic Drive Systems does not publish a clean customer-name breakout, so the mix is best read by end-market on a reporting basis. Industrial robot manufacturers — the FANUC, Yaskawa, ABB, KUKA tier and their Chinese equivalents — are the largest bucket at roughly 45% of revenue. Semiconductor and flat-panel-display manufacturing equipment is roughly 25%, and it is the most volatile slice: it was the swing factor in the FY2025 profit collapse. Aerospace, optical, medical and automotive together make up the remaining ~30%.

Humanoid-specific robotics is still a small slice — realistically under 5% of FY2025 revenue — but it is the fastest-growing one and the entire reason the multiple is where it is. The structural shift is the story: in 2023 humanoid was a rounding error and the investment case was “GDP-plus industrial automation compounder.” By early 2026, with the MinebeaMitsumi micro-actuator co-development, the 62% capex step-up, and management explicitly attributing order upside to robot makers, the mix is tilting. Goldman Sachs Japan’s framing — that each humanoid needs dozens of precision actuators and that this could expand Harmonic Drive Systems’ addressable market five-to-ten-fold over a decade — is the consensus bull narrative. Treat the “5-10x TAM” as an estimate, not a forecast; what is confirmed is that the order book is inflecting and the company is spending against it.

What’s actually happening at the humanoid OEMs

The uncomfortable truth sits at Tesla. Tesla's Optimus is the single largest potential demand pool for compact reducers, and Tesla has *not* standardised on Harmonic Drive Systems. Per multiple supply-chain reports, Tesla's primary harmonic-reducer supplier for Optimus is China's Suzhou Green Harmonic, which prices 30-40% below the Japanese incumbent and is reportedly targeting roughly 60% of the Optimus reducer market by 2026, with a new Suzhou plant ramping to 500,000 units of annual capacity. Suzhou Green Harmonic alone is cited at roughly 25% global harmonic-drive share. If the highest-volume Western humanoid program is being supplied predominantly by a Chinese low-cost vendor, the bull case for Harmonic Drive Systems cannot rest on Optimus.

Where Harmonic Drive Systems *does* win is the precision tier: programs where joint accuracy, lifetime and zero-backlash repeatability cannot be compromised, and the higher-mix Western and Japanese humanoid developers (Figure, Agility, Apptronik, and the Japanese cobot ecosystem) plus the entire installed base of industrial-robot OEMs who already qualified Harmonic Drive Systems decades ago. The MinebeaMitsumi micro-actuator partnership is the company's attempt to push *down* into the fingertip-count volume tier with a cost-engineered product rather than ceding it entirely. The qualification cycles here are long — a humanoid OEM that designs in a reducer keeps it for years — so the 2026-2027 design-win season is the period that sets share for the volume ramp. That is the mechanism to watch.

The competitive threat / Leaderdrive and Suzhou Green Harmonic

The competitive threat is the most concrete bear case on this name and it is not subtle. Chinese strain-wave makers — Leaderdrive (Shenzhen-listed), Suzhou Green Harmonic, and Zhejiang Laifual — have moved from under 5% of units installed in Chinese-assembled robots in 2018 to more than 35% by 2024, per industry data, with Leaderdrive taking the majority of that domestic share and roughly 15% of the *global* market by 2023. Domestic Chinese harmonic reducers are priced at roughly 40-60% of equivalent Harmonic Drive Systems units. Laifual launched miniature harmonic drives specifically for humanoid robotics in December 2024. This is a textbook oligopoly-erosion-by-cost-competition story, and it is happening in real time.

The counter-argument — and it is a real one, not a hopeful one — is that Chinese reducers still trail on precision-tier metrics: backlash consistency over lifetime, torque ripple, and the reliability data that a Western humanoid OEM needs before committing to a multi-year production design. Harmonic Drive Systems' moat is the qualification base and the metallurgy, not the patent (the core patents expired long ago). But "the Japanese

product is better" is exactly what the Japanese machine-tool and the Japanese solar industries said before they lost. The honest read: Harmonic Drive Systems keeps the precision tier and loses the volume tier on price, and the investment question is whether the precision tier grows fast enough in absolute yen to overwhelm the share loss. That is a genuinely open question.

The terminal risk

The terminal risk is not a technology transition away from strain-wave gears — humanoid robots will need compact rotary reducers for the foreseeable future, and the strain-wave architecture is hard to displace for distal joints. The terminal risk is *margin*, not *volume*: that the humanoid era arrives, Harmonic Drive Systems' unit volumes genuinely 5-10x, and yet the company earns a commodity-component margin on it because Chinese capacity has set the global clearing price 40-60% below the historical Japanese level. In that world the stock is a low-margin volume compounder dressed up as a monopoly, and the 85x trough multiple was a trap. The secondary terminal risk is architectural: some humanoid designers (and the broader 2026 industry literature) are pushing toward integrated linear actuators and planetary roller screws for high-force hip and knee joints, which would route incremental humanoid actuator dollars to THK and the screw makers rather than to strain-wave. Harmonic Drive Systems' credible defence is the MinebeaMitsumi micro-actuator (cost-engineered, hand-targeted) and the medium-term capex — but the defence is unproven until FY03/27 margins print.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. The single most-levered listed pure-play on humanoid joint count. Every credible humanoid architecture uses compact strain-wave reducers in the upper body, and Harmonic Drive Systems is the brand-defining incumbent. Goldman Sachs Japan's five-to-ten-fold TAM-expansion framing — driven by dozens of precision actuators per robot — is the structural case, and unlike most "robotics exposure" names, there is no segment dilution here: this is what the company does.

2. You are buying the cyclical trough, not the peak. FY2025 operating profit collapsed 94.4% to ¥1,967 million on a semiconductor-equipment correction with SG&A held up for the ramp. The medium-term plan targets ¥90 billion revenue and ¥15 billion

operating profit by FY03/27 — roughly 8x the FY2025 operating profit. The 85x forward P/E is a trough multiple; on FY03/27 plan earnings it compresses dramatically.

3. Order momentum is already inflecting. 9-month FY2026 sales rose 4.5% to ¥42.18 billion with management explicitly stating orders to robot makers exceeded forecast assumptions, across Japan, China and Europe. The semiconductor-equipment business is also recovering. The recovery is in the run-rate, not just the slideware.

4. The MinebeaMitsumi micro-actuator is a credible volume-tier play. The jointly-developed 13×19.4×60.4mm geared micro actuator, demonstrated at CES January 2026, targets humanoid hands and fingers — the highest unit-count part of the body. Co-developing with MinebeaMitsumi's manufacturing scale is the right answer to the cost problem, and it shows management understands the threat.

5. Capex is being spent against the thesis. FY2026 capital investment guided up 62% year-over-year to ¥6,100 million, explicitly for reducer capacity. Management is not merely talking the humanoid story; it is committing balance sheet to it ahead of the volume.

Gap

1. Chinese cost competition is eroding the oligopoly in real time. Domestic Chinese harmonic makers went from under 5% of Chinese-robot installs in 2018 to over 35% by 2024; Leaderdrive holds roughly 15% global share; Chinese units price 40-60% below Harmonic Drive Systems. This is not a future risk — it is a-tense share loss, and it caps the achievable margin.

2. Tesla Optimus — the largest single demand pool — is not a Harmonic Drive Systems program. Suzhou Green Harmonic is Optimus's primary harmonic-reducer supplier, targeting ~60% of that market by 2026 at 30-40% below Japanese pricing, with a 500,000-unit Suzhou plant ramping. The biggest Western humanoid volume story is largely closed to this name.

3. The valuation requires the recovery to be flawless. At 85x forward earnings on a trough base, any slippage in the FY03/27 ¥15 billion operating-profit target — a semiconductor double-dip, a slower humanoid design-win season, yen strength compressing export margins — leaves the multiple badly exposed. There is no valuation cushion.

4. The architecture could route humanoid dollars elsewhere. The 2026 industry shift toward integrated linear actuators and planetary roller screws for high-force hip/knee joints would direct incremental humanoid actuator spend to THK and screw makers, not to strain-wave. Harmonic Drive Systems owns the upper-body joint, but the upper body may not be where the BOM dollars concentrate.

Optionality

Event	Date / window	Direction
FY2026 (FYE Mar 2026) full-year results	~May 19, 2026	Binary on margin recovery vs FY03/27 plan
MinebeaMitsumi micro-actuator commercial orders	H2 2026	Bull if a named humanoid OEM designs it in
FY03/27 medium-term plan delivery (¥90B/¥15B)	Through Mar 2027	Binary — the whole valuation rests here
Humanoid OEM 2026-2027 design-win season	2026-2027	Bull/Bear — sets share for the volume ramp
Leaderdrive / Suzhou Green Harmonic global share data	Ongoing	Bear if share crosses into precision tier

---

The trade

Harmonic Drive Systems is the cleanest listed expression of the humanoid-joint thesis and simultaneously the name most exposed to the Chinese cost attack — which is precisely why it is a Bucket B, not Bucket A. Initiate at JPY 6,110-6,750 (current JPY 6,430 ±5%, RSI 69.1 is momentum in an early-cycle theme, not a sell signal), size at 1.5% of risk capital given the 85x trough multiple leaves no error margin, and stop at JPY 5,300 (below the 50-day moving-average structure and the recent consolidation floor). The defining binary is the FY2026 full-year print around May 19, 2026 — the number that shows whether robot-order momentum is converting to margin against the ¥15 billion FY03/27 target. If the thesis is right but you want a less single-point-of-failure expression, MinebeaMitsumi (6479) gives you the same humanoid micro-actuator optionality bolted onto a ¥1.6 trillion diversified bearings-and-motors base trading at 18x — a far cheaper way to own the same partnership, with the joint-layer concentration risk diversified away. **Conviction: 7 / 10.**

Sources referenced inline throughout: Harmonic Drive Systems FY2025 results briefing (May 20, 2025) and 9M FY2026 disclosure; Goldman Sachs Japan robotics framing; CES 2026 MinebeaMitsumi co-development announcement; Tesla Optimus supply-chain reporting (Suzhou Green Harmonic); industry data on Leaderdrive / Laifual Chinese

strain-wave share. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

6479 — MinebeaMitsumi / · BUY (Tier-1) · Conv 7/10 · Bucket B

MinebeaMitsumi (6479)

The miniature-ball-bearing world leader, with a Harmonic-Drive-co-developed humanoid micro-actuator — humanoid optionality at a conglomerate's 18x multiple.

Investment Research · Photoncap-style deep dive · v1 of "MinebeaMitsumi" · May 14, 2026

What MinebeaMitsumi physically does

MinebeaMitsumi is, at its foundation, the world's leading maker of miniature and small-size ball bearings — the tiny, precision-ground bearings, often only a few millimetres across, that let a shaft spin smoothly inside a motor, a fan, a hard-disk drive, a gimbal, or a robot joint. A bearing is the most basic enabling component in any rotating machine: without a low-friction, precisely-toleranced bearing, a motor cannot spin efficiently and a joint cannot articulate cleanly. MinebeaMitsumi grinds these to sub-micron tolerances at enormous volume, and that miniature-bearing franchise is the company's highest-margin, most-defensible core.

But MinebeaMitsumi is not a pure bearing company — it is an eight-business precision conglomerate, assembled through the 2017 Minebea-Mitsumi merger and subsequent acquisitions. Alongside bearings it makes small precision motors (stepping motors, brushless DC motors, fan motors), electronic devices and components, semiconductors (analog and power), optical devices, sensing devices, machined components, and access products. For the robotics theme, the relevant pieces are the bearings, the small precision motors, the frameless motors, the ball screws, the position sensors, and — increasingly — integrated actuators that combine these into a humanoid-ready module.

The reason MinebeaMitsumi belongs in the joint-layer batch is that a humanoid robot is, mechanically, a dense collection of small rotating joints and small motors, and every one of them needs miniature bearings, many need small precision motors, and the hands need micro-actuators. MinebeaMitsumi supplies the *enabling components* across

the body rather than owning one signature joint architecture. NMB, its components arm, explicitly markets bearings, ball screws, frameless motors, position sensors and cooling fans for humanoid robot systems. The company does not make the strain-wave gear (Harmonic Drive Systems) or the RV gear (Nabtesco) as a standalone product — but, critically, it co-develops with Harmonic Drive Systems, which is the single most important fact in this name's humanoid case.

Product roadmap

The component catalogue spans deep-groove and angular-contact miniature ball bearings, the hybrid and PM stepping motor lines, brushless DC and fan motor families, frameless torque motors, ball screws, and a broad sensing portfolio. MinebeaMitsumi's strategic direction is "INTEGappPRO" — the stated push to combine its eight businesses' components into higher-value integrated modules rather than selling discrete parts.

The roadmap event that defines this name's humanoid thesis is the high-torque micro actuator co-developed with Harmonic Drive Systems, given its first public demonstration at CES 2026 (January 2026). MinebeaMitsumi made its CES debut specifically to unveil humanoid robotics solutions, and the centrepiece was a robot hand employing geared micro actuators measuring 13mm wide, 19.4mm high and 60.4mm long, co-developed with Harmonic Drive Systems, shown alongside the related bearings and components. This is a genuinely important product positioning: it pairs Harmonic Drive Systems' strain-wave gear expertise with MinebeaMitsumi's miniature-motor and mass-manufacturing scale, aimed at the robot *hand and fingers* — the highest-actuator-count, highest-unit-volume part of a humanoid. NMB also markets a broader humanoid component set: bearings, ball screws, semiconductors and frameless motors. What MinebeaMitsumi does not have is a single proprietary humanoid joint it owns end-to-end — its humanoid roadmap is "components plus the Harmonic Drive partnership," not "our own reducer."

The financial print

MinebeaMitsumi runs a March fiscal year. The fiscal year ended March 2025 closed with net sales up 8.6% year-over-year to ¥1,522.7 billion (roughly $10 billion), operating income up 28.5% to ¥94,482 million, and profit attributable to owners up 10.0% to ¥59,457 million — a clean, double-digit-profit-growth year. For the fiscal year ending March 2026, the nine-month figures (April-December 2025, reported February 2026) showed cumulative sales rising and operating income of ¥75,208 million, up 3.1% year-

over-year, with the third quarter alone delivering ¥453.9 billion of sales (up 22.8%) and ¥30.8 billion of operating profit (up 17.8%). The company raised its full-year FY-March-2026 operating-income guidance to ¥101 billion on net sales of ¥1.60 trillion — above its prior ¥1.55 trillion estimate — with net income guided around ¥71 billion, factoring in roughly ¥4 billion of restructuring costs.

The defining feature of the financial print is the valuation gap. At a JPY 3,978 share price and roughly JPY 1.6 trillion market cap, the forward P/E of 18.5 is *less than a quarter* of Harmonic Drive Systems' 85x and a third of THK's and Nabtesco's multiples. MinebeaMitsumi trades like the diversified industrial conglomerate it is — because most of its earnings come from mature electronics, motors and semiconductors, not from humanoid robotics. That is the entire investment proposition: you are not paying a humanoid premium, so the humanoid micro-actuator is a near-free option. The next binary is the FY-March-2026 full-year result, expected on or around **May 13, 2026** — the print that confirms whether the ¥101 billion operating-income guidance held and whether management quantifies any humanoid traction.

Customer mix today

MinebeaMitsumi's customer mix is best read by business segment, because the conglomerate structure means no single customer or end-market dominates. The Precision Components segment — miniature ball bearings, machined components, rod ends — is the margin engine and the most defensible franchise, selling into motors, hard drives, aerospace, automotive and industrial customers globally. The Electronic Devices & Components segment (motors, sensing, backlights, switches, electronic devices) is the largest by revenue and includes the Mitsumi-heritage businesses serving smartphones, consumer electronics and automotive. The Semiconductors, Optical and Access businesses round out the eight-segment structure.

Humanoid-specific revenue is, honestly, tiny — it does not register as a disclosed slice, and MinebeaMitsumi's humanoid exposure is entirely an emerging-optionality story carried by the Harmonic Drive co-developed micro actuator and the NMB component set. The structural-shift framing here is different from the pure-play gear names: MinebeaMitsumi is not transitioning *from* an old business *to* humanoid; it is a diversified compounder that has *added* a credible humanoid call option without changing its valuation. The 2024 case was "diversified Japanese precision-components and electronics compounder at a low-teens-to-high-teens multiple." The 2026 case is the same, plus a co-developed humanoid micro-actuator demonstrated at CES — and the multiple has barely moved. That is the asymmetry: the downside is a cheap, profitable, diversified industrial; the upside is whatever the micro-actuator partnership converts

to.

What’s actually happening at the Harmonic Drive partnership

The mechanism that gives MinebeaMitsumi a real humanoid claim — rather than a generic “we sell components” claim — is the co-development with Harmonic Drive Systems. The logic of the partnership is sound: Harmonic Drive Systems owns the strain-wave gear know-how but is a relatively small company (¥608 billion market cap) without mass-consumer-electronics manufacturing scale; MinebeaMitsumi owns miniature-motor design, sub-micron bearing grinding, and the high-volume, cost-engineered manufacturing base built over decades of making bearings and motors by the billion. Pairing the two targets exactly the humanoid bill-of-materials problem that has worried this whole sector — that Japanese precision is too expensive for the volume tier. The 13×19.4×60.4mm geared micro actuator, demonstrated in a working robot hand at CES January 2026, is the first concrete output, and it is aimed at the highest-volume part of the humanoid body.

What is *not* yet happening is conversion to a named, dated production order. The CES demonstration is a capability showcase, not a design win, and neither company has disclosed a humanoid OEM committing the micro actuator to production. The 2026-2027 design-win season is the period to watch — and because the micro actuator targets robot hands, where actuator counts are highest, a single design win at a serious humanoid program would be a meaningful volume event for both partners. The honest read: MinebeaMitsumi has the right partner, the right manufacturing scale, and a demonstrated product, and it is being given essentially zero credit for any of it at 18x. That is the trade.

The competitive threat / Chinese miniature-bearing makers and the conglomerate discount

MinebeaMitsumi’s competitive threat is two-layered. On the legacy base, Chinese miniature-bearing and small-motor makers have spent years building scale and pricing below the Japanese incumbents — the same oligopoly-erosion pattern as in the gear names, applied to bearings and motors. MinebeaMitsumi’s defence is its sub-micron grinding tolerances, its scale, and the breadth of qualified customer relationships, but the legacy electronics-and-motors businesses are mature, competitive and not where the value-creation story lives. On the humanoid micro-actuator specifically, the compet-

itive set is the entire Chinese actuator supply chain ramping into Optimus and the other humanoid programs — the same Sanhua, Suzhou Green Harmonic and broader Chinese actuator ecosystem that is contesting the rotary-gear names. MinebeaMitsumi's edge there is precisely the Harmonic Drive partnership and its own manufacturing-cost discipline; whether that is enough to win volume-tier humanoid hand sockets against Chinese pricing is the open question.

But the more immediate "competitive threat" for the *stock* is internal: the conglomerate discount. MinebeaMitsumi trades at 18x because the market values it as a sum of mostly-mature parts, and any humanoid upside risks being diluted to invisibility inside a ¥1.6 trillion, eight-business revenue base. Even a successful micro-actuator program would need to scale enormously before it moved the consolidated needle. That is the structural reason this is a Bucket B and not a Bucket A — the option is cheap, but the conglomerate wrapper means it may take years for any humanoid success to show up in the multiple.

The terminal risk

MinebeaMitsumi's terminal risk is the inverse of Harmonic Drive Systems' — not "the humanoid era arrives but the margin is commoditised," but "the humanoid era arrives and it simply does not matter to a ¥1.6 trillion diversified industrial." If the micro-actuator partnership succeeds modestly, it is a rounding error against the electronics, motors and semiconductor base, and MinebeaMitsumi remains a low-teens-multiple conglomerate that was never really a robotics stock. The secondary terminal risk is the conventional one for any conglomerate: that the mature legacy businesses — consumer-electronics motors, backlights, machined components — face structural margin pressure from Chinese competition and slow consumer-hardware end-markets, dragging the whole group's earnings even as the bearing core and the humanoid option do well. MinebeaMitsumi's credible defence is the INTEGappPRO integration strategy (turning components into higher-value modules) and the genuinely strong, defensible miniature-bearing franchise — but the honest framing is that the bull case here is "cheap option, limited downside," not "this re-rates to a robotics multiple." Anyone underwriting a robotics-multiple re-rating is misreading the structure.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. Humanoid optionality at no premium. MinebeaMitsumi trades at 18.5x forward earnings — under a quarter of Harmonic Drive Systems' 85x, a third of THK's and Nabtesco's. The market prices it as a diversified conglomerate, so the co-developed humanoid micro actuator is close to a free option. That asymmetry is the whole trade.

2. The Harmonic Drive Systems co-development is a genuinely well-matched partnership. Pairing Harmonic Drive Systems' strain-wave know-how with MinebeaMitsumi's miniature-motor design and decades of high-volume, cost-engineered manufacturing directly attacks the humanoid BOM-cost problem. The 13×19.4×60.4mm micro actuator, demonstrated in a working robot hand at CES January 2026, targets the highest-volume part of the body.

3. The base business is compounding and guidance was raised. FY-March-2025 delivered net sales up 8.6% to ¥1,522.7 billion and operating income up 28.5% to ¥94.5 billion; FY-March-2026 guidance was raised to ¥101 billion operating income on ¥1.60 trillion sales. Q3 alone grew sales 22.8%. This is a profitable, growing industrial even before any humanoid credit.

4. The miniature-bearing core is a real, defensible franchise. World leadership in sub-micron-tolerance miniature ball bearings is a durable moat — bearings are the most basic enabling component in every rotating machine, and MinebeaMitsumi's grinding scale and customer qualifications are hard to replicate. This is the high-quality floor under the conglomerate.

5. Limited downside. Unlike the 85x-trough Harmonic Drive Systems or the 84-RSI THK, the worst realistic case for MinebeaMitsumi is "cheap, profitable, diversified industrial that was never a robotics stock" — the 18x multiple is its own valuation cushion.

Gap

1. The conglomerate discount may never lift. At 18x with an eight-business, ¥1.6 trillion revenue base, even a successful micro-actuator program would be diluted to near-invisibility on the consolidated P&L. The option is cheap precisely because the structure makes it hard for humanoid success to ever show up in the multiple.

2. No named humanoid production design win. The CES January 2026 demonstration is a capability showcase, not an order. Neither MinebeaMitsumi nor Harmonic Drive Systems has disclosed a humanoid OEM committing the micro actuator to production — the humanoid case is still entirely forward-looking.

3. The chart is overbought. RSI 84.4 and +35.3% versus the 50-day moving average is a stretched entry — second only to THK in this batch. Even in an early-cycle theme, an

84-RSI print raises near-term mean-reversion risk and argues for buying on a pullback rather than chasing.

4. The legacy businesses face Chinese competition and soft consumer-hardware markets. The mature electronics-and-motors segments — consumer-device motors, backlights, machined components — are exposed to Chinese pricing and weak consumer hardware demand, which could drag group earnings even as the bearing core and humanoid option do well.

Optionality

Event	Date / window	Direction
FY-March-2026 full-year results	~May 13, 2026	Binary on ¥101B guidance + humanoid commentary
First named micro-actuator production design win	2026-2027	Bull — would force a partial re-rating
Q1 FY-March-2027 results	~August 7, 2026	Binary on base-business momentum
INTEGappPRO module-integration progress	Through 2026-2027	Bull if it lifts the conglomerate multiple
Chinese miniature-bearing / actuator share trends	Ongoing	Bear if legacy-base margins compress

---

The trade

MinebeaMitsumi is the value-side way to own the Japanese joint-layer humanoid thesis: a profitable, growing, diversified ¥1.6 trillion industrial with a genuinely well-matched Harmonic Drive Systems micro-actuator partnership — and a multiple that gives you essentially none of that option for free. The catch, and the reason it is Bucket B rather than A, is the conglomerate discount: humanoid success may take years to surface against the eight-business base. With RSI at 84.4 the entry needs discipline — initiate on a pullback into JPY 3,779-4,177 (current JPY 3,978 ±5%) and favour the lower half of the band, size at 1.5% of risk capital, and stop at JPY 3,350 (below the 50-day structure and the pre-breakout base). The defining binary is the FY-March-2026 full-year print

around May 13, 2026 — the test of the ¥101 billion operating-income guidance and any humanoid quantification. If you want maximum humanoid torque from the *same partnership* and will pay the trough multiple for it, the pivot is Harmonic Drive Systems (6324) — the pure-play strain-wave name; MinebeaMitsumi is the inverse trade — own the partnership's manufacturing-scale half at a conglomerate price, with the joint-layer concentration risk diversified away and a hard valuation floor underneath. **Conviction: 7 / 10.**

Sources referenced inline throughout: MinebeaMitsumi FY-March-2025 results (released May 9, 2025) and 9-month FY-March-2026 disclosure (February 2026); MinebeaMitsumi FY-March-2026 guidance raise; CES 2026 humanoid robot hand and Harmonic Drive Systems co-development announcement (January 2026); NMB Technologies humanoid component materials. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

SHA — Schaeffler AG · BUY (Tier-1) · Conv 7/10 · Bucket B

Schaeffler AG (SHA)

A cheap German bearings-and-powertrain giant with the most concrete humanoid actuator timeline in the joint-layer batch — series production starts Q2 2026.

Investment Research · Photoncap-style deep dive · v1 of "Schaeffler" · May 14, 2026

— Native listing: Schaeffler AG, Xetra ticker SHA (yfinance: SHA0.DE) —

What Schaeffler physically does

Schaeffler is one of the world's great precision-bearing companies — a German Mittelstand-grown industrial giant whose core competence is the rolling-element bearing and everything adjacent to it: the metallurgy, the grinding, the tolerancing, and the system engineering that turns a bearing into a complete motion module. A bearing, again, is the foundational enabling component of any rotating machine — and Schaeffler makes them across the entire size range, from tiny precision bearings to the multi-metre bearings inside wind-turbine nacelles. Around that core it has built large automotive and industrial businesses: e-mobility drive systems and components for electric vehicles, conventional powertrain and chassis systems, an enormous automotive aftermarket business, and the Bearings & Industrial Solutions division serving wind, aerospace, rail, and general industry.

For the robotics theme, Schaeffler's claim is the humanoid joint *actuator* — and it is a credible one, because a humanoid actuator is, mechanically, exactly the kind of integrated motion module Schaeffler already builds for cars and machines: a bearing, a gearbox, a motor interface, packaged for a specific load case. Schaeffler is developing rotary actuators, compact gearboxes and linear actuators for humanoid joints, sitting alongside its precision bearings as the enabling components of a robot's articulation. It does not own a signature reducer architecture the way Harmonic Drive Systems (strain-wave) or Nabtesco (RV) do — but it brings something none of the Japanese pure-plays brings: a €23.5 billion industrial balance sheet, a German-engineering brand, and the

manufacturing footprint to ramp an actuator platform across regions at automotive scale and cost.

Since 2025, Schaeffler has organised these growth ambitions into a dedicated “Others” division, combining humanoid robotics and defence — explicitly separating the future-oriented bets from the core so investors can see them. The stated strategic goal is to generate up to 10% of group revenue from these new growth areas by 2035.

Product roadmap

The robotics roadmap centrepiece is Schaeffler’s rotary actuator platform — a cost-competitive family spanning sizes XXS to XL that the company says covers roughly 80% of humanoid-joint market demand. The platform won the Hermes Award (an industrial-technology prize awarded at Hannover Messe), a credible third-party validation, and — critically for a Photoncap dating discipline — Schaeffler has given it a concrete commercialisation timeline: *first series production is planned for the second quarter of 2026*, with ramp-ups scheduled across all relevant regions through the remainder of the year. Alongside the rotary actuator platform, Schaeffler is developing compact gearboxes and linear actuators for humanoid joints, and its precision bearings feed every joint regardless of actuator type.

Just as important as the product is the customer roadmap. Schaeffler has announced partnerships with four named humanoid manufacturers — Agility Robotics, Neura Robotics, Humanoid, and Leju Robotics — plus the German defence company Helsing on the defence side. Four named humanoid customers and a dated Q2 2026 series-production start is, bluntly, the most concrete humanoid commercialisation evidence of any name in this joint-layer batch: where Harmonic Drive Systems has a CES demonstration and THK has a product shelf with no named win, Schaeffler has named OEM partnerships and a production date on the calendar. The honest caveat: “partnership” and “series production planned” are not the same as disclosed production volumes or revenue — but the specificity here is real and unusual.

The financial print

Schaeffler reported full-year 2025 results on March 5, 2026. Group revenue was €23,492 million, down 0.6% on a pro-forma, constant-currency basis from €24,313 million — essentially flat, in a tough automotive market. EBIT before special items improved to €936 million, an EBIT margin before special items of 4.0%, up from a

pro-forma 3.5% the prior year. Free cash flow before M&A was positive €266 million, a sharp improvement from a pro-forma negative €694 million. By division: Bearings & Industrial Solutions — the most relevant segment for the robotics thesis — grew revenue 0.7% to €6,368 million with an EBIT margin before special items of 7.5% (€475 million), helped by wind and aerospace bearings; Vehicle Lifetime Solutions (the aftermarket) is the highest-margin division. The board proposed a dividend of €0.30 per share.

For 2026, Schaeffler guides group revenue of roughly €22.5-24.5 billion (constant-currency growth of -4.3% to +4.3%) and an EBIT margin before special items of 3.5-5.5%. The Bearings & Industrial Solutions division is guided to a 7.0-9.0% EBIT margin on roughly €6.2-6.7 billion of revenue. At a €9.91 share price and roughly €9.4 billion market cap, the forward P/E of 14.2 is a low-multiple, automotive-supplier-style valuation — Schaeffler is *not* being priced as a robotics stock, which is the core of the asymmetry: the humanoid actuator business is being valued at close to nothing inside a cheap industrial. The next binary is the Q2 2026 results, expected around **August 5, 2026** — but the more thesis-specific catalyst is the Q2 2026 first series production of the rotary actuator platform itself, the moment "planned" becomes "shipping."

Customer mix today

Schaeffler's revenue mix is dominated by automotive. The E-Mobility and Powertrain & Chassis divisions — components and systems for electric and conventional vehicles — together make up the bulk of the €23.5 billion group. Vehicle Lifetime Solutions, the automotive aftermarket business, is roughly €3.1-3.3 billion and is the highest-margin division. Bearings & Industrial Solutions is roughly €6.2-6.7 billion, serving wind, aerospace, rail and general industry — and it is the division where the robotics actuator and humanoid-bearing revenue will land. The new "Others" division houses the humanoid-robotics and defence growth bets.

Humanoid-specific revenue today is tiny — it sits inside the "Others" growth division precisely because it is not yet a material revenue contributor; the 10%-of-revenue target is a *2035* aspiration, not a current figure. The structural-shift story is genuine but slow: the 2024 Schaeffler case was "low-multiple German automotive-and-industrial supplier navigating the EV transition and a soft car market." The 2026 case adds a credible, dated, named-customer humanoid actuator business — but it adds it to a €23.5 billion automotive-heavy base that will, for years, dwarf it. The honest framing matches MinebeaMitsumi's: this is humanoid optionality bolted onto a cheap diversified industrial, not a humanoid pure-play. The difference from MinebeaMitsumi is that Schaeffler's humanoid evidence — four named OEMs, a dated production start — is more concrete,

while its core (automotive) is arguably more challenged than MinebeaMitsumi's (precision components and electronics).

What's actually happening at the named humanoid customers

The mechanism that makes Schaeffler's humanoid claim more than a press release is the specificity of the four named partnerships — Agility Robotics, Neura Robotics, Humanoid, and Leju Robotics — paired with the Q2 2026 series-production date. Agility Robotics (the maker of the Digit warehouse humanoid) and Neura Robotics (a well-funded German humanoid and cognitive-robotics developer) are serious, commercially-oriented programs; Humanoid and Leju Robotics extend the reach. The logic of why these OEMs would choose Schaeffler is sound: a humanoid developer wants an actuator partner who can ramp to volume at automotive cost and quality, with multi-region manufacturing — and that is precisely Schaeffler's structural advantage over the smaller Japanese pure-plays. The XXS-to-XL rotary actuator platform covering ~80% of joint demand is designed exactly to be the broad-line, qualify-once supplier across a humanoid's body.

What is *not* yet disclosed is the conversion to volume: Schaeffler has not published production unit commitments or humanoid revenue figures, and "partnership" can mean anything from a co-development MOU to a committed supply agreement. The Q2 2026 first-series-production milestone is the moment to watch — if Schaeffler ships actuators in series to a named OEM on schedule, the option moves from credible to real. The 2026 ramp-up "across all relevant regions" is the company's stated plan; whether it executes on time, and whether the named OEMs themselves hit their own production ramps, is the open question. But on the evidence available in mid-2026, Schaeffler has done more to de-risk the "will it actually commercialise" question than anyone else in this batch.

The competitive threat / the joint-layer field and Chinese cost

Schaeffler's competitive set in humanoid actuators is the entire joint layer plus the Chinese supply chain. Against the Japanese pure-plays — Harmonic Drive Systems, Nabtesco — Schaeffler does not have a signature reducer architecture, so where a humanoid joint specifically needs a strain-wave or RV gear, Schaeffler is at a disadvantage; its play is the *integrated rotary actuator as a module*, competing on system integration, breadth, cost and ramp capability rather than on owning the gear. Against the Chinese actuator and reducer ecosystem — Suzhou Green Harmonic, Sanhua, the broad

low-cost field supplying Tesla Optimus — Schaeffler faces the same pricing pressure every Western and Japanese precision incumbent faces, and one of the named customers, Leju Robotics, is itself Chinese, which cuts both ways (access to the Chinese humanoid market, but also exposure to Chinese cost expectations).

Schaeffler's genuine competitive edge is the one the Japanese pure-plays lack: scale and cost engineering. The rotary actuator platform is explicitly positioned as *cost-competitive*, designed by a company that has spent a century engineering automotive components down to a price. If the humanoid-actuator market commoditises — and the whole joint-layer analysis suggests it will — Schaeffler is structurally better placed than Harmonic Drive Systems or Nabtesco to compete in the commoditised volume tier, because automotive-grade cost discipline is its native habitat. The risk is the inverse: if the market stays a precision-premium market, Schaeffler lacks the proprietary gear moat to capture that premium. Schaeffler is the cost-competitive entrant, not the precision-monopoly incumbent.

The terminal risk

Schaeffler's terminal risk is dilution by the core. Even if the humanoid actuator business executes perfectly — Q2 2026 series production on time, all four named OEMs ramping, the platform winning broad joint share — the 2035 target is for new growth areas (humanoid *and* defence combined) to reach just ~10% of group revenue. That means for the next several years, Schaeffler's earnings and share price will be set overwhelmingly by the ~€20 billion automotive-heavy core: the EV transition, the health of the global car market, European auto-supplier margins, and the pace of the conventional-powertrain decline. A humanoid success can be entirely real and still be invisible against an automotive cyclical downturn. The secondary terminal risk is execution: a young actuator platform from a company without a humanoid track record could slip its Q2 2026 date, hit ramp problems, or find its named OEM customers miss their own production timelines — humanoid OEM schedules have a poor punctuality record industry-wide. Schaeffler's credible defence is the low entry multiple (14.2x means you are not paying for humanoid success) and the cost-engineering edge — but the honest framing, as with MinebeaMitsumi, is "cheap option with a hard valuation floor," not "this re-rates to a robotics multiple soon."

Bull / Gap / Optionality (Photoncap framing)

Bull

1. The most concrete humanoid commercialisation timeline in the joint-layer batch. Schaeffler has a dated Q2 2026 first-series-production milestone for its rotary actuator platform, four named humanoid OEM partnerships (Agility Robotics, Neura Robotics, Humanoid, Leju Robotics), and a Hermes Award for the platform. No other name in this batch pairs named customers with a production date.

2. Humanoid optionality at an automotive-supplier multiple. At 14.2x forward earnings and ~€9.4 billion market cap, Schaeffler is priced as a cheap, soft-market auto supplier — the humanoid actuator business is valued at close to nothing. Like MinebeaMitsumi, the asymmetry is a near-free option with a hard valuation floor.

3. Scale and cost engineering are the right weapons for a commoditising market. The rotary actuator platform is explicitly cost-competitive, spanning XXS-XL to cover ~80% of joint demand, built by a century-old automotive-cost-engineering company. If humanoid actuators commoditise — the base-case for the whole joint layer — Schaeffler is structurally better placed than the Japanese precision pure-plays to win the volume tier.

4. The core is stabilising and cash-generative. 2025 EBIT margin before special items improved to 4.0% from 3.5%, free cash flow swung to +€266 million from -€694 million pro-forma, and Bearings & Industrial Solutions — the robotics-relevant division — runs a 7.5% EBIT margin with wind and aerospace tailwinds. This is not a distressed core.

5. A dedicated "Others" division gives investors visibility. By ring-fencing humanoid robotics and defence in a separate reporting division since 2025, Schaeffler lets the market see the growth bets discretely rather than burying them — which makes an eventual re-rating more achievable if the numbers come through.

Gap

1. The humanoid business is dilutively small for years. The 10%-of-revenue target for new growth areas is a 2035 figure, and it bundles humanoid *with* defence. For the foreseeable future, Schaeffler's earnings are set by a ~€20 billion automotive-heavy core, not by actuators.

2. The core is automotive — cyclical and mid-EV-transition. Group 2026 guidance allows for revenue down 4.3%, and the EBIT margin guide of 3.5-5.5% is thin. A weak car market or a rocky EV transition can swamp any humanoid progress entirely.

3. No proprietary reducer moat. Unlike Harmonic Drive Systems (strain-wave) or Nabtesco (RV), Schaeffler does not own a signature gear architecture. Its play is the

integrated module on cost and breadth — strong if the market commoditises, weak if it stays a precision-premium market.

4. Execution and customer-timing risk on an unproven platform. A first-generation actuator platform from a company with no humanoid track record could slip its Q2 2026 date or hit ramp issues, and the four named OEM customers could miss their own production timelines — humanoid OEM schedules are chronically optimistic.

Optionality

Event	Date / window	Direction
First series production of the rotary actuator platform	Q2 2026	Bull — “planned” becomes “shipping”
Q2 2026 results	~August 5, 2026	Binary on core margin + actuator-ramp commentary
Disclosed humanoid production volumes / revenue	2026-2027	Bull — would force a partial re-rating
Named-OEM (Agility / Neura) production ramps	2026-2027	Bull/Bear — Schaeffler’s volume depends on theirs
2026 EU automotive market + EV-transition trajectory	Through 2026	Bull/Bear — sets the core earnings

---

The trade

Schaeffler is the Western mirror of MinebeaMitsumi: humanoid actuator optionality bolted onto a cheap, diversified industrial trading at 14.2x — but with a *more concrete* humanoid evidence base (four named OEMs, a dated Q2 2026 production start) and a *more challenged* core (European automotive mid-EV-transition versus MinebeaMitsumi’s precision components and electronics). That balance — best-in-batch commercialisation specificity, weakest-in-batch core — earns it Bucket B, conviction 6. Initiate at €9.41-10.41 (current €9.91 ±5%; RSI 73.0 is momentum in an early-cycle theme, and the +29.7% vs 50MA is moderate by this batch’s standards), size at 1.0% of risk capital given the automotive-cyclical drag on the core, and stop at €8.30 (below the 50-day structure and the pre-breakout base). The defining catalyst is the Q2 2026 first series produc-

tion of the rotary actuator platform — the moment the most-credible humanoid timeline in the batch either delivers or slips — with the Q2 2026 earnings print around August 5 as the financial binary. If you want the same "cheap industrial, free humanoid option" structure with a sounder core and a Harmonic Drive Systems partnership instead of an in-house platform, MinebeaMitsumi (6479) is the cleaner expression; Schaeffler is the trade for the investor who wants the most tangible, dated, named-customer humanoid commercialisation evidence available today and will accept the European auto-cycle baggage to get it. **Conviction: 6 / 10.**

Sources referenced inline throughout: Schaeffler full-year 2025 results press release (March 5, 2026); Schaeffler Q1 2026 disclosure on the rotary actuator platform and humanoid robotics push; Schaeffler announcements of partnerships with Agility Robotics, Neura Robotics, Humanoid, Leju Robotics and Helsing; Schaeffler 2026 divisional guidance; Hermes Award for the rotary actuator platform. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

6481 — THK / THK · WATCH (Tier-2) · Conv 6/10 · Bucket C

THK / THK (6481)

The linear-motion-guide leader, slimmed to a pure industrial-machinery play — a real but unproven roller-screw angle on humanoid legs, into a badly over-bought tape.

Investment Research · Photoncap-style deep dive · v1 of "THK" · May 14, 2026

What THK physically does

THK invented the linear motion guide — the LM guide — and the company name is an abbreviation of *Toughness, High quality, Know-how*. Where Harmonic Drive Systems and Nabtesco solve rotary motion, THK solves *straight-line* motion: how to move a heavy carriage along a rail with near-zero friction, high rigidity, and micron-level repeatability. An LM guide is a rail with precision-ground grooves and a block that rides on it via recirculating ball bearings; it is the part that lets a machine-tool table, a semiconductor wafer stage, or a 3D printer head travel precisely. THK also makes ball screws and ball splines (which convert rotary motor motion into linear thrust), LM guide actuators (the rail, screw, motor and bearing integrated into a unit), and electric and linear actuators. It is, in short, the precision-linear-motion company.

The reason this matters for the robotics theme is the humanoid leg. A humanoid robot's hip and knee joints carry the highest forces in the body — they support and propel the entire mass — and the 2026 engineering consensus is increasingly that those joints are best served not by a rotary reducer but by a *linear* actuator built around a planetary roller screw. A roller screw replaces a ball screw's point-contact balls with line-contact threaded rollers, so it carries several times the load in the same envelope and lasts far longer under the cyclic shock of walking. That is precisely THK's domain. THK does not make the strain-wave wrist gear (Harmonic Drive Systems) or the RV base gear (Nabtesco); THK's humanoid claim is the legs and the linear high-force joints, plus a separate angle in robot hands and grippers.

The corporate structure changed materially in 2025-2026, and it changes how to read

the company. THK divested its Automotive & Transportation business, transforming itself from a two-segment company into a focused Industrial Machinery (IMS) pure-play. From the fiscal year ending December 2025 onward, the reported numbers exclude automotive entirely — so historical comparisons need care, but the resulting entity is a cleaner, higher-mix precision-motion business.

Product roadmap

The core catalogue is the LM Guide family — the SHS, SSR, HSR and SRG series spanning the size and rigidity range for machine tools, semiconductor equipment and general industrial machinery — plus the ball screw and ball spline lines and the KR/SKR LM guide actuator series that integrate motion into a ready-to-mount unit. THK also runs an FA Solutions push, layering IoT, condition-monitoring and AI-driven predictive maintenance onto the mechanical base, as part of a stated strategy to evolve from a pure component maker into a services-and-modules company.

The roadmap items that matter for this theme are the SEED Solutions actuator family and the robot-hand work. THK's SEED actuators are compact, modular electric actuators that pair with the SEED Driver — a motor driver only 20mm square — to give humanoid and service-robot integrators a small, self-contained motion building block; SEED actuators enable simple ON/OFF gripping and pinching regardless of workpiece shape. THK also developed the General-Purpose Gripper TRX, a highly compact three-finger universal gripper with a single motor, a link mechanism and an integrated controller. On the high-force side, THK's planetary roller screws are the product that maps directly to humanoid hip and knee actuators, where roller-screw load capacity and life are the differentiators against ball screws. The honest framing: THK has the *right products* for the humanoid leg and hand, but it has not yet announced a named, dated humanoid production design win — the roadmap is product-ready, not order-confirmed.

The financial print

THK switched to a December fiscal year-end, so FY2025 is the fiscal year ended December 2025; the FY2025 4Q investor information was released February 12, 2026, and the FY2026 1Q investor information was released May 11, 2026. Because the Automotive & Transportation business was divested, both the FY2025 actuals and the FY2026 forecast exclude automotive — so the headline numbers look smaller than the old two-segment THK. For FY2026 (year ending December 2026), THK projects revenue around ¥276

billion and operating income around ¥31 billion on the slimmed industrial-machinery base.

The genuinely positive signal is the upgrade. THK sharply raised its FY2026 first-half and full-year forecasts, citing stronger-than-expected orders in the industrial-machinery business in Japan and overseas: it now expects first-half consolidated revenue of ¥138 billion and operating income of ¥15.4 billion — gains of 8.7% and 51.0% respectively over its previous outlook. A 51% operating-income guidance raise is a real inflection, not a rounding adjustment. The counter-signal is the recent history: FY2025 second-quarter operating profit fell 29% year-over-year to ¥4.5 billion, so THK is climbing out of a soft patch rather than compounding from strength. At a JPY 7,675 share price and roughly JPY 860 billion market cap, the forward P/E of 39.2 is rich for a linear-motion maker and embeds a humanoid premium that the order book does not yet substantiate. The Q1 FY2026 print already landed on May 11, 2026; the next binary is the Q2 FY2026 result, expected around **August 7, 2026** — the test of whether the raised guidance holds.

Customer mix today

THK does not publish a named-customer table; the mix reads by end-market within the now-singular Industrial Machinery segment. Machine tools are a major bucket — THK LM guides are spec'd into the machine-tool builders' platforms — alongside general industrial machinery, semiconductor and flat-panel-display manufacturing equipment, electronic-device assembly equipment, and industrial robots. The semiconductor and FPD equipment exposure is the swing factor: it drove the 2024-2025 softness and is part of the recovery now lifting guidance. With automotive gone, the residual mix is higher-margin and more clearly tied to the capex cycles of precision manufacturing.

Humanoid-specific revenue is, candidly, negligible today — it does not register as a reported slice. THK's humanoid exposure is entirely forward-looking, carried by the SEED actuators, the TRX gripper, and the planetary roller screws, none of which has a disclosed humanoid production design win. The structural-shift story THK wants investors to believe is that the 2024 case ("cyclical machine-tool and semiconductor-equipment recovery") becomes, through 2026-2028, "the same recovery *plus* the linear high-force humanoid leg." That is plausible — the engineering logic for roller-screw legs is sound — but it is the least-proven humanoid call in this batch. Where Harmonic Drive Systems has a co-developed micro actuator demonstrated at CES and Nabtesco has launched named compact-gear families with a quantified share estimate, THK has the right component shelf and no named customer. The honest read: you are buying an industrial-machinery recovery at 39x with a humanoid option attached that is real in engineering

terms and speculative in revenue terms.

What’s actually happening at the humanoid leg

The substantive question for THK is whether humanoid legs go linear, and if so, whose roller screws win. The 2026 engineering literature is increasingly clear on the first half: integrated linear motion is displacing traditional rotary systems for high-force joints, and inverted planetary roller screws deliver a far better power-to-weight ratio for the hip and knee, where distributed load contact lets a roller screw support several times a ball screw’s load while maintaining long operational life. Tesla’s Optimus reference architecture uses 14 linear actuators alongside its 14 rotary ones — the linear count is not a niche; it is half the body’s actuators. So the demand pool THK is aiming at is genuine and large.

The mechanism that is *not* yet working in THK’s favour is the design-win conversion. The high-volume Western humanoid program — Optimus — has routed much of its actuator supply chain to China, and the roller-screw and lead-screw content for the one-million-robot Optimus plan is being heavily sourced from Chinese makers. THK competes in the precision tier, where its roller-screw life and accuracy data is a genuine differentiator, and it has the SEED actuator system for integrators who want a packaged module. But the qualification cycle is the gate: a humanoid OEM that designs in a roller-screw supplier keeps it for the production life, and THK has not announced winning one of those sockets in production. The 2026-2027 design-win season is decisive — and until a named win prints, the humanoid leg story is an engineering thesis, not a revenue forecast.

The competitive threat / Hiwin and the Chinese roller-screw makers

THK’s competitive set on the linear side is led by Taiwan’s Hiwin — a credible, lower-cost, scaled LM-guide and ball-screw maker that has spent years taking share from THK in the price-sensitive industrial tier — alongside the Chinese roller-screw and lead-screw makers ramping capacity into the Optimus supply chain. The dynamic is the same oligopoly-erosion pattern seen at the rotary-gear names: a Japanese precision incumbent with the quality lead and the installed-base relationships, facing Taiwanese and Chinese challengers who price well below it and improve fast. On the humanoid leg specifically, the threat is acute precisely because it is a greenfield market — there is

no THK installed base to defend, so the Chinese and Taiwanese makers compete for the roller-screw sockets from the first design cycle, on price, with no incumbency inertia working for THK.

The mitigant is that roller screws — particularly inverted planetary roller screws — are a genuinely demanding precision product, harder to manufacture to humanoid-grade life specifications than a commodity ball screw, and THK's metallurgy and grinding know-how is a real edge in that specific sub-component. But "real edge in a sub-component" is exactly the position Harmonic Drive Systems and Nabtesco are in, and it has not stopped Chinese share gains there. The realistic outcome is consistent across the whole joint layer: THK keeps the precision-tier roller-screw business and competes hard on price for the volume tier, and the question is absolute-yen growth versus share loss.

The terminal risk

THK's terminal risk is the cleanest "commoditisation before scale" case in this batch. The bull thesis requires humanoid legs to (1) go linear, (2) standardise on roller screws, and (3) source those roller screws meaningfully from THK rather than from Hiwin or Chinese makers. Each step is plausible; the conjunction is not assured. If humanoid legs instead use a rotary architecture, the linear-leg demand pool shrinks and the dollars route to the rotary-gear names. If legs do go linear but the roller-screw socket is won on price by lower-cost competitors — the most likely failure mode given the greenfield-market dynamic — THK is left as a strong industrial-machinery recovery story with a disappointingly thin humanoid attach, and the 39x multiple was overpaying for an option that did not pay off. THK's credible defences are the roller-screw manufacturing edge, the SEED actuator system as a packaged module, and the FA Solutions services layer — but none of these is yet validated by a named humanoid production order, which is why the terminal risk weighs heavier here than at the two rotary-gear leaders.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. The undisputed linear-motion-guide leader, now a clean industrial-machinery pure-play. THK invented the LM guide and leads the global market. The 2025-2026 automotive divestiture removed the most cyclical, lowest-mix drag, leaving a focused, higher-margin Industrial Machinery entity — a structurally better business to own through the cycle.

2. Guidance was just raised hard. THK lifted its FY2026 first-half operating-income forecast 51.0% and revenue 8.7% on stronger-than-expected industrial-machinery orders in Japan and overseas — first-half revenue now ¥138 billion, operating income ¥15.4 billion. A 51% guidance raise is a real inflection in the order book.

3. THK owns the right products for the humanoid leg. Planetary roller screws are the 2026-consensus answer for high-force hip and knee actuators, and roller screws are core THK territory. Optimus's reference architecture alone uses 14 linear actuators — the linear demand pool is half the body, not a niche.

4. The SEED actuator system and TRX gripper give THK a packaged-module angle. SEED actuators plus the 20mm-square SEED Driver, and the single-motor three-finger TRX universal gripper, let THK sell integrators a self-contained motion building block rather than a bare component — capturing more BOM value if a design win lands.

5. The semiconductor and machine-tool capex recovery is a real near-term earnings leg. The same precision-manufacturing capex cycle lifting Nabtesco is lifting THK; the guidance raise is evidence the recovery is in the order book, not just the slideware — so the base business is improving even before any humanoid revenue.

Gap

1. The tape is badly overbought. RSI 84.8 and +41.3% versus the 50-day moving average is the most stretched setup in this batch. Even in an early-cycle theme where high RSI is momentum rather than an automatic sell, an 84.8 print signals a poor entry — the risk of a sharp 10-15% mean-reversion before any thesis development is high.

2. The humanoid story has no named production design win. Unlike Harmonic Drive Systems (CES-demonstrated co-developed micro actuator) or Nabtesco (launched compact-gear families with a share estimate), THK has the component shelf and zero disclosed humanoid production orders. The 39x multiple is paying for an unconfirmed option.

3. The humanoid leg socket is a greenfield fight THK can lose on price. With no installed base to defend, THK competes for roller-screw legs against Hiwin and Chinese makers from the first design cycle, on price, with no incumbency inertia. Greenfield markets are exactly where Japanese precision incumbents have historically lost share fastest.

4. The recent earnings base is soft, not strong. FY2025 Q2 operating profit fell 29% year-over-year to ¥4.5 billion. THK is climbing out of a soft patch on a guidance raise, not compounding from strength — if the industrial-machinery recovery stalls, the 39x multiple has no cushion and no proven humanoid revenue to fall back on.

Optionality

Event	Date / window	Direction
Q2 FY2026 results	~August 7, 2026	Binary on the raised guidance holding
First named humanoid roller-screw / SEED design win	2026-2027	Bull — would re-rate the humanoid option
Humanoid leg architecture: linear vs rotary consensus	2026-2027	Bull/Bear — determines if THK's pool exists
Machine-tool / semiconductor capex cycle trajectory	Through 2026	Bull/Bear — drives the near-term earnings
Hiwin / Chinese roller-screw share in humanoid supply chain	Ongoing	Bear if THK is shut out of volume sockets

---

The trade

THK has the right shelf of products for the humanoid leg and a genuine industrial-machinery recovery in hand — but it carries the least-proven humanoid revenue path in this batch and, critically, the most overbought chart, which together drop it to Bucket C. The discipline call here is to *not chase*: with RSI at 84.8 and the stock +41.3% over the 50-day moving average, initiate only on a pullback into JPY 7,291-8,059 (current JPY 7,675 ±5%) and ideally toward the lower half of that band, size light at 1.0% of risk capital given the unconfirmed humanoid option and the stretched tape, and stop at JPY 6,300 (below the 50-day structure and the pre-breakout base). The Q1 FY2026 print already landed May 11; the next binary is the Q2 result around August 7, 2026 — the test of whether the 51% guidance raise holds. If you want the linear-motion-into-humanoid thesis with a confirmed humanoid partnership and a far cheaper multiple, the pivot is MinebeaMitsumi (6479) — its 18x multiple and ¥1.6 trillion diversified base make it the lower-risk way to own Japanese precision-motion humanoid optionality, leaving THK as the higher-beta, buy-the-dip-only expression. **Conviction: 6 / 10.**

Sources referenced inline throughout: THK FY2025 4Q investor information (February 12, 2026) and FY2026 1Q investor information (May 11, 2026); THK FY2026 guidance-

raise disclosure; THK SEED Solutions and General-Purpose Gripper TRX product materials; 2026 humanoid-actuator engineering literature on planetary roller screws. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

NOVT — Novanta Inc. · WATCH (Tier-2) · Conv 6/10 · Bucket B

Novanta Inc. (NOVT)

A high-quality precision-motion, photonics and robotic-tooling compounder riding AI-and-automation tailwinds — but a robotics name only at the edges, priced at 37x.

Investment Research · Photoncap-style deep dive · v1 of "Novanta" · May 14, 2026

What Novanta physically does

Novanta is a precision-technology compounder — a company that has spent two decades acquiring and integrating a portfolio of high-margin, engineering-intensive niche businesses in photonics, precision motion and medical technology. It is not a humanoid-joint pure-play, and it should not be read as one. What Novanta is, for the robotics theme, is a supplier of the *enabling subsystems* that make precision motion and machine vision possible — encoders, motors, motion-control mechatronics, laser beam-steering systems, and robot end-of-arm tooling — across both industrial and medical robotics.

Three pieces of the portfolio matter here. First, Celera Motion: an award-winning provider of precision motion-control components and subsystems — optical and inductive encoders, frameless and direct-drive motors, and customised mechatronic modules. An encoder is the sensor that tells a control system exactly where a joint is; it is the feedback element of every closed-loop robot joint, and Celera Motion is a genuine leader in the precision-encoder niche. Second, ATI Industrial Automation: described as the world's leading engineering-based developer of robotic accessories and robot-arm tooling — automatic tool changers, multi-axis force/torque sensors, utility couplers, material-removal and compliance tools. ATI is the end-of-arm-tooling layer that sits between an industrial robot's wrist and whatever it is manipulating. Third, the photonics businesses: laser sources and the intelligent laser beam-steering and scanning subsystems used in advanced manufacturing, semiconductor and medical applications.

So Novanta's robotics exposure is real but specific: it is the encoder, the motion mechatronic, the force/torque sensor and the tool changer — the feedback-and-interface layer

of robotics — rather than the actuator, the reducer or the bearing that this joint-layer batch is principally about. It belongs in the batch as the “precision motion and sensing” adjacency, not as a strain-wave or RV competitor.

Product roadmap

Novanta’s robotics-and-automation product stack runs, in its own framing, from servo controllers to sensors to motion components. On the Celera Motion side, the catalogue is the precision encoder families (optical and inductive), the frameless and direct-drive motor lines, and customised mechatronic subsystems for medical and industrial robotics. On the ATI Industrial Automation side, it is the automatic tool changers, the multi-axis force/torque sensing systems, the utility couplers, and the material-removal and compliance tooling — the comprehensive end-of-arm-tooling portfolio. On the photonics side, the roadmap event Novanta highlighted on its Q1 CY2026 call is the newly launched intelligent laser beam-steering subsystems, with proprietary capabilities aimed at probe-card production for AI GPU chips, laser additive manufacturing for aerospace and drones, advanced packaging and substrate production for data-center applications, and light engines for deep-UV and EUV lithography.

The roadmap discipline point: Novanta’s product cadence is strong and measurable — new-product sales grew over 50% year-over-year in Q1 CY2026, lifting the “Vitality Index” (share of revenue from recently launched products) to 27% of sales, and company-wide design wins were up nearly 30% year-over-year. That is a healthy innovation engine. But there is no humanoid-specific product line with a launch date — Novanta’s robotics roadmap is industrial-and-medical-robotics-and-AI-manufacturing, and humanoid is, at most, a future incremental application for its encoders and motion components rather than a named program. Read Novanta as a precision-motion-and-photonics compounder with a robotics tilt, not as a humanoid roadmap story.

The financial print

Novanta reported full-year CY2025 results on February 23, 2026: GAAP revenue up 3% to $981 million, GAAP diluted EPS of $1.47, adjusted diluted EPS of $3.29, and adjusted EBITDA of $221 million. By segment, Automation Enabling Technologies generated revenue of $500.8 million at a 47.8% gross margin and $114.5 million operating profit, while Medical Solutions delivered $479.8 million at a 41.6% gross margin and $51.2 million operating profit — two roughly equal-sized segments, with Automation the higher-margin

one. During the year Novanta acquired Keonn Technologies for $75.1 million (RFID, into Medical Solutions), issued $613.1 million net via tangible equity units, and put in place a new $1.0 billion credit facility maturing June 2030 — an active-balance-sheet year consistent with the acquisitive compounder model.

The Q1 CY2026 print, released around May 11-12, 2026, was a clear beat: revenue of $257.7 million, up 10.4% year-over-year (10% reported, 3% organic), non-GAAP EPS of $0.81, 4.3% above consensus, adjusted EBITDA up 14% with margin up 70 basis points, and — the standout — bookings up 37% year-over-year at a 1.1 book-to-bill, with double-digit bookings growth in every business. Novanta raised full-year CY2026 guidance to revenue of $1.04-1.055 billion and adjusted diluted EPS of $3.50-3.65. At the $154.66 reference price and roughly $5.5 billion market cap, the forward P/E of 37.1 is a premium-compounder multiple — it reflects the genuine quality and the AI-automation tailwind, but it leaves little room for disappointment. The next binary is the Q2 CY2026 earnings, expected around **August 4, 2026** — the test of whether the 37% bookings surge sustains into revenue.

Customer mix today

Novanta discloses its mix by end-market and by segment. By end-market, the medical markets accounted for roughly 53% of CY2025 revenue and advanced industrial markets roughly 47% — a deliberate, balanced split. By segment, Automation Enabling Technologies (~$501 million) houses the precision motion, encoders, robotic tooling, laser beam-steering and bearing-spindle businesses, serving advanced industrial processes, industrial and medical robotics, and medical lasers; Medical Solutions (~$480 million) houses the medical-grade technologies plus its own robotics-and-automation content. The customer base is broad and diversified across surgical-device OEMs, semiconductor and advanced-manufacturing equipment makers, and industrial-automation integrators — Novanta does not have dangerous single-customer concentration.

Humanoid-specific revenue is not a disclosed slice and is, candidly, immaterial today — Novanta's robotics exposure is overwhelmingly *surgical and industrial* robotics, not humanoid. The structural-shift story Novanta tells is "AI-driven robotics and automation, minimally invasive and robotic surgery, digital manufacturing, precision medicine" — all genuine secular tailwinds, but none of them is the humanoid joint-count story that drives Harmonic Drive Systems or Nabtesco. The honest framing for this batch: Novanta is the highest-business-quality name here and simultaneously the one with the *least* direct humanoid leverage. An investor including Novanta in a robotics theme is buying a diversified precision-tech compounder with a robotics-and-AI tilt — if humanoid robots scale, Novanta's encoders and motion components get an incremental tailwind, but the thesis

does not stand or fall on humanoids the way the pure-play joint names' theses do.

What's actually happening at the AI-and-automation demand pool

The mechanism driving Novanta right now is not humanoid joint count — it is the broader AI-and-automation capex wave, and it is showing up clearly in the numbers. Q1 CY2026 bookings up 37% year-over-year with a 1.1 book-to-bill, double-digit bookings growth in *every* business, design wins up nearly 30%, and new-product sales up over 50% — that is a company whose end-markets are inflecting across the board. The specific drivers Novanta cites are concrete: probe-card production for AI GPU chips (its laser beam-steering subsystems), advanced packaging and substrate production for data centers, laser additive manufacturing for aerospace and drones, light engines for deep-UV and EUV lithography, and the rising automation and digitization of manufacturing lines demanding higher throughput, smaller form factors and tighter tolerances. On the medical side, the robotic-surgery and minimally-invasive-procedure tailwind drives the encoder and motion-component content.

What is happening, in short, is that Novanta is a high-quality picks-and-shovels supplier to *multiple* secular waves at once — AI semiconductor manufacturing, data-center buildout, surgical robotics, industrial automation — and several of them are accelerating together. That breadth is the bull case and the bear case simultaneously: it makes Novanta resilient and genuinely growthy, but it also means the humanoid-robotics angle is one tailwind among many, not the defining one. For a joint-layer batch focused on humanoid leverage, the honest mechanism read is that Novanta is the least-humanoid-levered name and the most-diversified-quality name — and which of those framings matters more depends on what the investor wants the position to do.

The competitive threat / the niche-by-niche field

Novanta does not face one named competitor the way Harmonic Drive Systems faces Leaderdrive — it faces a different competitive set in each niche. In precision encoders, Celera Motion competes against Renishaw, Heidenhain and other precision-metrology specialists. In robotic end-of-arm tooling, ATI Industrial Automation competes against Schunk, OnRobot, Zimmer and a fragmented gripper-and-tool-changer field — though ATI's position as a leading engineering-based developer in tool changers and force/torque sensing is genuinely strong. In photonics and laser beam-steering,

Novanta competes against Coherent, II-VI-heritage businesses, and various scanning-optics specialists. The pattern is that Novanta is a strong number-one or number-two in several defensible niches rather than a dominant monopolist in one — which is a more resilient structure than a single-product pure-play, but also means no single moat carries the whole company.

The Chinese-cost-competition theme that dominates the rotary-gear names is less acute for Novanta, because its products are higher up the engineering-and-IP stack — precision encoders, force/torque sensors and laser beam-steering subsystems are harder to commoditise than a bearing or a gear, and Novanta's medical exposure (53% of revenue) sits behind regulatory qualification barriers that slow competitive entry. The genuine competitive risk is subtler: in a multi-niche compounder, any one niche losing share to a sharper specialist is survivable, but a pattern of niche-by-niche erosion would slowly degrade the premium multiple. There is no active IP litigation flagged here; the competitive threat is the ordinary one of staying ahead, niche by niche, in a portfolio of competitive markets.

The terminal risk

Novanta's terminal risk is not commoditisation or a technology cliff — its diversification and its position up the IP stack protect it from the single-point failures that threaten the gear pure-plays. The terminal risk is *valuation versus identity*: Novanta trades at 37x forward earnings as a premium AI-and-robotics compounder, and if the market eventually re-frames it as what it more precisely is — a good, diversified, mid-single-digit-organic-growth precision-instruments company with a cyclical advanced-industrial half — the multiple has a long way to compress. The 37x is underwritten by the current bookings surge and the AI-manufacturing tailwind; a normalisation of that tailwind, or a couple of quarters where the 37% bookings growth does not convert to revenue, would expose the multiple. The secondary risk is the acquisitive-compounder risk: Novanta's model depends on continuing to buy and integrate niche businesses accretively (Keonn, the Schneider Motion and ATI deals before it), and a bad acquisition or an integration stumble would dent both earnings and the premium. For *this batch specifically*, the relevant "terminal" framing is narrower: Novanta is the name where humanoid robotics is least load-bearing — so the risk is simply that an investor who bought it *for* humanoid exposure finds the humanoid contribution never becomes visible, while the actual drivers (surgical robotics, AI semis, data center) are ones they could have bought more directly elsewhere.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. The highest business quality in the joint-layer batch. Two roughly equal $480-500M segments, 41-48% gross margins, a diversified medical/industrial split, leadership positions in multiple defensible niches (Celera Motion encoders, ATI robotic tooling, laser beam-steering). This is a genuinely high-quality compounder, not a single-product bet.

2. The bookings inflection is broad and real. Q1 CY2026 bookings up 37% year-over-year at 1.1 book-to-bill, with double-digit bookings growth in *every* business, design wins up ~30%, and new-product sales up over 50% (Vitality Index at 27%). That is an across-the-board demand acceleration, and management raised full-year guidance on it.

3. Multiple secular tailwinds at once. AI GPU probe-card production, data-center advanced packaging, EUV light engines, surgical robotics, industrial automation — Novanta is a picks-and-shovels supplier to several accelerating waves simultaneously, which makes the growth durable rather than dependent on one theme.

4. Up the IP stack, insulated from Chinese-cost commoditisation. Precision encoders, force/torque sensors and laser beam-steering subsystems are harder to commoditise than gears or bearings, and the 53% medical revenue sits behind regulatory qualification barriers — Novanta is largely outside the oligopoly-erosion dynamic that threatens the rotary-gear names.

5. A proven, disciplined acquisitive compounder. The Keonn, Schneider Motion and ATI deals, the $613M equity-unit raise and the new $1.0B credit facility show a management team that has executed the buy-and-integrate model accretively over many years.

Gap

1. The multiple leaves no room for disappointment. 37.1x forward earnings prices Novanta as a premium AI-robotics compounder. If the bookings surge does not convert, or the AI-manufacturing tailwind normalises, the multiple has substantial downside — and the stock is already +21.7% over its 50-day moving average with RSI 68.9.

2. It is the least humanoid-levered name in the batch. Novanta's robotics exposure is surgical and industrial robotics plus AI-manufacturing — humanoid is, at most, a minor future incremental tailwind. An investor buying Novanta *for* the humanoid theme is buying the wrong vehicle.

3. Multi-niche means no single dominant moat. Novanta is a strong #1 or #2 in several niches against sharp specialists (Renishaw, Heidenhain, Schunk, Coherent). Niche-by-niche share erosion is survivable individually but would slowly degrade the premium multiple if it became a pattern.

4. Compounder risk: the model depends on continued accretive M&A. A bad acquisition or an integration stumble would hit both earnings and the premium — and the balance sheet has been actively used (equity units, $1.0B facility), so the model is leveraged to keep working.

Optionality

Event	Date / window	Direction
Q2 CY2026 earnings	~August 4, 2026	Binary — does the 37% bookings surge convert to revenue?
AI semiconductor / data-center capex trajectory	Through 2026	Bull/Bear — a primary current demand driver
Surgical-robotics OEM program ramps	2026-2027	Bull — drives Medical Solutions encoder content
Next bolt-on acquisition	2026	Bull if accretive and well-integrated
Any humanoid-robotics encoder/tooling design win	2026-2027	Bull — would add a tailwind it currently lacks

---

The trade

Novanta is the quality anchor of this joint-layer batch and, paradoxically, the name least dependent on the humanoid thesis that defines the batch — a high-margin, diversified precision-motion-and-photonics compounder with a genuine across-the-board bookings inflection, riding AI-manufacturing, data-center and surgical-robotics tailwinds. That quality, plus the broad demand acceleration, earns it Bucket B; the 37x multiple and the thin direct humanoid leverage cap conviction at 6. Initiate at $146.93-162.39 (current $154.66 ±5%; RSI 68.9 is momentum in an early-cycle theme, not an automatic sell, but the multiple argues for the lower half of the band), size at 1.0% of risk capital given the premium valuation leaves no error margin, and stop at $132.00 (below the 50-day

structure and the pre-breakout consolidation). The defining binary is the Q2 CY2026 print around August 4 — the test of whether the 37% bookings growth converts to revenue and justifies the raised guidance. The honest pivot for *this theme*: if the investor's goal is direct humanoid joint-layer leverage, Novanta is not the cleanest expression — Harmonic Drive Systems (6324) or Schaeffler (SHA) carry the actual humanoid-actuator torque; Novanta is the right holding instead for the investor who wants a high-quality, diversified precision-tech compounder that *participates* in robotics and AI automation broadly, with surgical robotics and AI-semis as the real engines and humanoid as a free upside option. **Conviction: 6 / 10.**

Sources referenced inline throughout: Novanta full-year CY2025 results (released February 23, 2026) and Q1 CY2026 results (released ~May 11-12, 2026); Novanta Q1 CY2026 earnings call commentary on bookings, Vitality Index and laser beam-steering subsystems; Novanta robotics-and-automation, Celera Motion and ATI Industrial Automation product materials; Novanta CY2026 raised guidance. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

ALNT — Allient Inc. · WATCH (Tier-2) · Conv 5/10 · Bucket C

Allient Inc. (ALNT)

A small-cap motion-control roll-up with a genuine frameless-torque-motor shelf and a loud humanoid push — but the humanoid revenue is still a whisper and Q1 just missed.

Investment Research · Photoncap-style deep dive · v1 of "Allient" · May 14, 2026

What Allient physically does

Allient is a small-cap American precision-motion company — the renamed and expanded successor to Allied Motion Technologies — that designs and manufactures motors, drives, and integrated motion-control systems. Its core competence is the electric motor and its control: brushless DC motors, frameless torque motors, slotless and coreless designs, plus the drives, controllers and integrated actuators that turn a motor into a usable motion subsystem. Like Nidec, Allient's relevance to robotics flows from the simple fact that the motor is the prime mover of every robot joint — but where Nidec is a $25-billion-revenue global behemoth, Allient is a roughly $1 billion market-cap, ~$555 million-revenue company assembled through a multi-year acquisition program (the "roll-up" that built Allied Motion into Allient).

The robotics — and specifically humanoid — claim rests on Allient's frameless and high-torque-density motor portfolio. A frameless motor is sold as just the rotor and stator, without a housing or bearings, so the robot designer can integrate it directly into the joint structure — exactly what a humanoid actuator wants, because it minimises mass and packaging. Allient's MegaFlux frameless direct-drive torque motors span twelve frame diameters from 60mm to 330mm with continuous stall torque up to 181 Nm, and its ThinGap motors are ultra-lightweight, zero-cogging designs aimed at aerospace and robotics. These are legitimate, well-regarded products for advanced robotics, gimbals and UAVs. What Allient does *not* have is a proprietary reducer architecture, a unique enabling technology, or scale — it is a competent broad-line motion-component supplier, and that framing matters for the whole thesis.

Allient has been explicit and vocal about the humanoid push. In April 2026 it published a whitepaper, "A Selection Guide to Motors for Humanoid Robotics Systems," examining how actuators enable locomotion, manipulation and balance, and it is showcasing its motion technologies at robotics and defence trade shows (XPONENTIAL and MDEX in May 2026). The intent is clear; the question is whether intent is translating to revenue.

Product roadmap

The robotics-relevant catalogue is led by the MegaFlux frameless direct-drive torque motor series — twelve frame diameters from 60mm to 330mm, multiple stack lengths per diameter, continuous stall torque up to 181 Nm — which Allient has been actively expanding (a line expansion was announced through 2025-2026). Alongside MegaFlux sit the ThinGap motors (ultra-lightweight, zero-cogging, aerospace-and-robotics-targeted), the KinetiMax brushless platform (compact, power-dense), and the SA Series axial-flux motors, which Allient expanded in June 2025 with a new 63.5mm outer-diameter variant — axial-flux being a topology well-suited to the flat, high-torque packaging humanoid joints favour. Allient also sells the drives, controllers and integrated actuator assemblies around these motors, positioning itself as a motion-subsystem supplier rather than a bare-component vendor.

The roadmap "events" that matter for this theme are softer than at the Japanese names: the April 2026 humanoid motor-selection whitepaper and the May 2026 trade-show showcases are marketing-and-positioning milestones, not product launches or design wins. Allient has not announced a named humanoid OEM customer or a dated humanoid production program. The honest Photoncap framing: Allient has a credible, expanding shelf of motors that *could* go into humanoid joints, and a management team actively marketing into the space — but there is no confirmed, dated humanoid design win to anchor a roadmap on. The roadmap is "we have the right motors and we are telling the market," not "we are shipping to a named program."

The financial print

Allient reported full-year CY2025 results on March 5-6, 2026: revenue of $554.5 million and net income of $22 million, with full-year gross margin of 32.8% — up 150 basis points year-over-year and a record for a full-year period — and adjusted EBITDA of $76.9 million (13.9% of revenue). Operating cash flow hit a record $56.7 million, up 35%, and the company deleveraged meaningfully, paying down roughly $48 million of

debt over the year. That was a genuinely solid year: margin expansion, record cash generation, balance-sheet repair.

Then the Q1 CY2026 print, released around May 6, 2026, was a stumble. Revenue rose 4.6% year-over-year to $138.9 million — in line with expectations — but non-GAAP EPS of $0.46 came in 13.6% below consensus, adjusted EBITDA of $17.28 million (12.4% margin) missed by 6.4%, operating margin fell to 6.7% from 7.7% a year earlier, and free cash flow margin dropped to 2.9% from 9.7%. The stock fell about 11.9% on the day, to roughly $68. Backlog was a relative bright spot at $251 million, up 5.8% year-over-year. At the $62.50 reference price and roughly $1.06 billion market cap, the forward P/E of 20.4 is moderate — but it sits on earnings that just missed and margins that just compressed. Note the chart context: unlike every overbought name in this batch, Allient has RSI 36.6 and is 6.8% *below* its 50-day moving average — it has already been sold off, which is both a risk signal (the market is worried) and a potential entry advantage (no chasing required). The next binary is the Q2 CY2026 earnings, expected around **August 6, 2026** — the print that tests whether Q1 was a one-quarter cost-and-mix issue or the start of a margin de-rating.

Customer mix today

Allient discloses its mix by end-market, and for full-year CY2025 it broke down as: Industrial roughly 48% of revenue (the largest bucket, up 8% on the year, driven by power-quality solutions into data-center infrastructure); Vehicle roughly 17% (down 6%, soft powersports partly offset by commercial automotive and construction); Medical roughly 15% (up 5% on steady surgical-instrument demand); Aerospace & Defense roughly 15% (down 5%, lumpy program timing, hurt by the M10 Booker tank program cancellation); and Distribution roughly 5%. It is a genuinely diversified small-cap industrial — no single end-market dominates, which provides resilience but also means no single secular tailwind drives it.

Humanoid-specific revenue is not a disclosed line item, and realistically it is a thin, unquantified slice buried inside the Industrial and Aerospace & Defense buckets. Allient's humanoid exposure is, at this stage, a *positioning* story — the whitepaper, the trade shows, the marketing of MegaFlux and ThinGap into the space — not a *revenue* story. The structural-shift framing is the weakest in this batch: there is no visible 2024-to-2026 mix change attributable to humanoid, because humanoid is not yet a measurable contributor. The bull would say Allient is early and the diversified base funds the wait; the skeptic would say a sub-scale roll-up loudly marketing into a hyped end-market, with no named win, is a familiar and often disappointing pattern. The honest read sits closer to the skeptic until a design win prints.

What’s actually happening at the humanoid opportunity

The mechanism question for Allient is uncomfortable: in a humanoid actuator market where the demand is being aggregated by a small number of large programs (Tesla Optimus, Figure, Agility, the Chinese humanoid OEMs), why would a serious humanoid developer choose a sub-scale American roll-up over a Schaeffler (scale, cost engineering, four named OEM partnerships already), a Nidec (the world’s largest motor maker, governance issues notwithstanding), a MinebeaMitsumi (manufacturing scale plus the Harmonic Drive partnership), or a low-cost Chinese motor supplier? Allient’s honest answer would be the niche: frameless and slotless motors for specific joints where its ThinGap zero-cogging or MegaFlux torque-density genuinely differentiates, or as a Western-sourced supplier for defence-adjacent or aerospace-grade humanoid programs where supply-chain provenance matters. That is a real but narrow lane.

What is *not* happening — at least not visibly as of May 2026 — is conversion. No named humanoid OEM, no disclosed humanoid production program, no quantified humanoid backlog. The April 2026 whitepaper and the May trade shows are Allient telling the market it wants to play; they are not evidence the market has chosen Allient. For a Photoncap analysis that demands “tool counts, qualification dates, competitor displacement specifics,” Allient simply does not yet have them in humanoid. The 2026-2027 design-win season is the period that will decide whether Allient’s humanoid push is a real niche business or marketing dressed on a diversified industrial — and an investor buying today is buying before that question is answered.

The competitive threat / scale and the joint-layer field

Allient’s competitive problem is structural: it is sub-scale in a field of giants. In humanoid motors specifically, it competes against Nidec (the world’s largest electric-motor maker), Schaeffler (€23.5 billion industrial with a dated actuator-production timeline and four named humanoid customers), MinebeaMitsumi (¥1.6 trillion, miniature-motor scale, Harmonic Drive partnership), and the entire Chinese motor and actuator supply chain that is winning Optimus volume on price. Against any of these, Allient’s roughly $555 million of total revenue and lack of a proprietary moat is a genuine disadvantage — it cannot out-invest, out-scale, or out-price the larger players, and it does not own a unique enabling technology the way the reducer pure-plays do.

Allient’s defensible position is the specialist niche: high-performance frameless and slot-

less motors (ThinGap, MegaFlux) for applications where torque density, low cogging or aerospace-grade quality matter more than unit cost, and Western/defence supply-chain provenance. That is a real business — it is roughly what Allient already does profitably across aerospace, defence and medical. But it is a *niche* humanoid play, not a volume-tier one, and the risk is that even the niche gets contested by larger players' specialist lines (Nidec, Schaeffler and others all have high-performance frameless offerings). There is no IP litigation to track here; the competitive threat is plain commercial — scale and breadth versus a sub-scale roll-up.

The terminal risk

Allient's terminal risk is that the humanoid attach never becomes material *and* the multiple gave it credit for it anyway. The bull case implicitly asks investors to pay something above a diversified-small-cap-industrial multiple for the humanoid optionality — but if humanoid revenue stays a thin niche slice, Allient is simply a competent, diversified, sub-scale motion-control roll-up worth a mid-teens industrial multiple, and the 20.4x forward P/E (on just-missed earnings) was an overpay. The secondary terminal risk is the one the Q1 print just flagged: margin. A roll-up's value depends on integrating acquisitions into expanding margins; Q1 CY2026 showed operating margin and free-cash-flow margin both compressing, and if that is a trend rather than a quarter, the entire roll-up thesis weakens regardless of humanoid. There is also concentration-of-bad-luck risk in the lumpy Aerospace & Defense book — the M10 Booker cancellation already cost it. Allient's credible defence is the diversified base (no single end-market collapse sinks it), the record cash generation and deleveraging in CY2025, and the genuinely good frameless-motor products — but none of that adds up to a humanoid *thesis*; it adds up to a fair, slightly-cheap small-cap industrial with a marketing campaign attached.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. A genuinely good, expanding frameless-motor product set. MegaFlux frameless direct-drive torque motors (twelve frame sizes, 60-330mm, up to 181 Nm) and ThinGap zero-cogging motors are well-regarded for advanced robotics, gimbals and UAVs — these are real products in the right category for humanoid joints, and Allient is actively expanding the lines.

2. The chart is the opposite of the rest of the batch — already sold off. RSI 36.6 and -6.8% versus the 50-day moving average, after an 11.9% post-earnings drop. There is no chasing risk here; the entry is into weakness, not a melt-up — a structurally better starting point than THK or Minebea.

3. CY2025 was a solid operational year. Full-year gross margin of 32.8% (a record, up 150bps), record operating cash flow of $56.7 million up 35%, and ~$48 million of debt paydown. The underlying business executed well before the Q1 stumble.

4. Diversified base funds the humanoid wait. Industrial ~48%, Vehicle ~17%, Medical ~15%, Aerospace & Defense ~15% — no single end-market dominates, so Allient can market into humanoid over multiple years without betting the company on it.

Gap

1. Q1 CY2026 missed on margin. Non-GAAP EPS 13.6% below consensus, adjusted EBITDA 6.4% short, operating margin down to 6.7% from 7.7%, free-cash-flow margin down to 2.9% from 9.7%. For a roll-up whose thesis is margin integration, that is the wrong signal — and the stock fell ~12% on it.

2. No named humanoid design win. The April 2026 whitepaper and May trade shows are positioning, not orders. There is no disclosed humanoid OEM customer, no dated humanoid production program, no quantified humanoid backlog. The humanoid case is entirely a marketing story today.

3. Sub-scale in a field of giants. Allient's ~$555 million of revenue and lack of a proprietary moat puts it against Nidec, Schaeffler, MinebeaMitsumi and the Chinese supply chain in humanoid motors. It cannot out-scale or out-price any of them; its only defensible lane is a narrow specialist niche.

4. The multiple pays for optionality the numbers do not yet support. At 20.4x forward earnings on just-missed results, Allient is priced above a plain diversified-small-cap-industrial multiple — implicitly crediting humanoid optionality that has produced zero disclosed revenue.

5. Lumpy Aerospace & Defense and soft Vehicle end-markets. The M10 Booker tank program cancellation already hit the A&D book (down 5% in CY2025), and Vehicle fell 6% on soft powersports — two of the four end-markets are currently headwinds.

Optionality

Event	Date / window	Direction
Q2 CY2026 earnings	~August 6, 2026	Binary — was Q1's margin miss a quarter or a trend?
First named humanoid design win	2026-2027	Bull — would validate the humanoid push
Backlog trajectory ($251M and growing?)	Quarterly	Bull/Bear — leading indicator for the diversified base
Aerospace & Defense program awards	2026	Bull if new programs replace the M10 Booker loss
Margin recovery toward CY2025's 13.9% adj. EBITDA	Through 2026	Bull — the roll-up thesis lives or dies here

---

The trade

Allient is the honest "interesting, not yet proven" name in this batch — a competent, diversified small-cap motion-control roll-up with genuinely good frameless motors and a loud humanoid marketing push, sold off after a Q1 margin miss, with no named humanoid win to underwrite the optionality. That combination — real products, real sell-off, no proof — earns it Bucket C, conviction 5. The one structural advantage is the entry: with RSI 36.6 and the stock below its 50-day moving average, you are buying weakness, not chasing strength, so the price-discipline framework actually favours initiating here. Initiate at $59.38-65.63 (current $62.50 ±5%), size *small* at 0.75% of risk capital — this is a speculative-attach position, not a core holding, given the unproven humanoid revenue and the just-flagged margin question — and stop at $52.00 (below the recent low and the structural support that the post-earnings drop established). The defining binary is the Q2 CY2026 print around August 6 — the test of whether Q1's margin compression was one quarter or a trend. If you want humanoid motor exposure with a real commercialisation timeline and named customers rather than a whitepaper, the cleaner expression is decisively Schaeffler (SHA) — same "cheap industrial, motor/actuator humanoid optionality" structure, but with four named OEM partnerships and a dated Q2 2026 production start, at a lower multiple; Allient is the trade only for the investor who specifically wants a small-cap, US-listed, defence-adjacent specialist motor name and will accept that the humanoid story is, for now, just a story. **Conviction: 5 / 10.**

Sources referenced inline throughout: Allient full-year CY2025 results (released March 5-6, 2026) and Q1 CY2026 earnings (released ~May 6, 2026); Allient "A Selection Guide to Motors for Humanoid Robotics Systems" whitepaper (April 2026); Allient MegaFlux and ThinGap product materials; Allient XPONENTIAL / MDEX 2026 trade-show announcements. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

6594 — Nidec / · SKIP / WAIT (Tier-3) · Conv 3/10 · Bucket D

Nidec (6594)

The world’s largest electric-motor maker — and, right now, an active accounting scandal you cannot underwrite. Capability is not the problem; credibility is.

Investment Research · Photoncap-style deep dive · v1 of “Nidec” · May 14, 2026

What Nidec physically does

Nidec is the largest electric-motor company on earth. It built that position on the brushless DC spindle motor that spins hard-disk drives, and then spent three decades acquiring and building its way into nearly every motor category that exists: automotive traction motors and e-axle drive systems for electric vehicles, appliance motors, commercial and industrial motors, small precision motors, fans, machine-tool and factory-automation motors, and motors for robotics. If a machine has a rotating shaft, Nidec almost certainly makes a motor for it somewhere in its sprawling portfolio.

For the robotics theme, the relevant fact is that the motor is the prime mover of every robot joint. A humanoid robot’s actuators are, at their heart, electric motors — usually compact, high-torque-density brushless or frameless designs — paired with a reducer and an encoder. Nidec has genuine, deep capability in exactly this kind of compact precision motor, and it has the manufacturing scale that the humanoid-cost problem demands. On pure engineering and manufacturing merit, Nidec belongs in any serious conversation about who supplies the humanoid motor.

But this deep dive cannot be written the way the others in this batch are written, because Nidec is, as of May 2026, in the middle of an active accounting scandal — and that fact dominates everything else. A company’s financial statements are the instrument through which an investor underwrites it. When those statements are under restatement, when results are delayed, and when the company itself is investigating whether its own directors bear legal liability, the normal apparatus of financial analysis does not apply. So this piece is structured honestly around that reality: the motor capability is real, and it is also, for now, beside the point.

Product roadmap

Nidec's product breadth is its defining feature: spindle motors, automotive traction motors and the E-Axle integrated EV drive system, appliance and HVAC motors, commercial and industrial motors, small precision motors, machine-tool spindle motors, and factory-automation and robotics motors. The E-Axle — Nidec's integrated motor-inverter-gearbox EV drive unit — was the centrepiece of the growth story for years and is, not coincidentally, at the heart of the accounting problem: the third-party committee's findings point to goodwill and fixed-asset impairments concentrated in the automotive business.

On robotics specifically, Nidec markets motors for factory automation and has the compact-precision-motor and frameless-motor capability that maps to humanoid actuators. But here the normal Photoncap discipline — name every product with a launch date, distinguish confirmed from estimate — runs into a wall: with the company's own financial reporting under investigation and results delayed, any roadmap claim must be treated as low-confidence. The honest framing is that Nidec *has the engineering* for humanoid motors and the manufacturing scale to make them at cost, but there is no clean, verifiable, dated humanoid product roadmap to anchor a thesis on, and even if there were, it would be subordinate to the governance situation.

The financial print

This is where the name breaks. Nidec has *delayed* the release of its results for the fiscal year ended March 31, 2026, because it is still working to correct prior-period financial statements and complete the current year's accounts. On February 27, 2026, a Third-Party Committee delivered an investigation report finding a range of accounting misconduct across the Nidec Group; a final supplementary report followed, and on April 17, 2026 the company announced its receipt and response. The misconduct described includes, among other items, failing to record inventory valuation losses by falsely recognising asset values in raw materials and products with extremely low future usage and sales prospects and no asset value — i.e., carrying worthless inventory at value. The restatements reach back to the fiscal year ended March 31, 2022.

The quantified damage so far: the scandal is expected to cost roughly ¥250 billion in impairment charges, primarily goodwill and fixed assets tied to the automotive business, with the company warning that additional impairment losses may follow as derivative

impacts of the downward revisions to prior years. Nidec has decided to forgo the year-end dividend for the fiscal year ending March 2026. It has established an Executive Responsibility Investigation Committee to determine whether current and former directors, auditors and executive officers bear legal liability. There is, as of mid-May 2026, no audited, reliable set of financial statements to value the company against — the forward P/E of 14.2 is computed off pre-restatement estimates that are, by the company's own admission, unreliable. The next genuine binary is the eventual release of the restated and delayed financial statements; the placeholder date of **June 30, 2026** is an estimate, not a confirmed date, and slippage is itself a bearish signal.

Customer mix today

Nidec's customer mix cannot be stated with the precision the house style demands, because the segment financials are under restatement. On a pre-scandal reporting basis, the automotive business — E-Axle traction systems and EV-related motors — had been the largest growth segment and is precisely the area where the impairments are concentrated, which tells you the segment's true profitability was materially worse than reported. The appliance, commercial and industrial motor businesses and the small-precision-motor business make up the diversified balance, and these are likely the more genuinely sound parts of the company. The hard-disk-drive spindle motor business is a mature, declining legacy.

Humanoid-specific revenue is a tiny emerging slice and, frankly, immaterial to the investment case here — the case is entirely about whether the company's financial reporting can be trusted, not about end-market mix. The structural shift that matters for Nidec is not "industrial to humanoid"; it is "from a company with credible accounts to a company without them, and back again, if it gets there." Until that round-trip completes, breaking out customer percentages would be false precision on numbers the company itself has disavowed.

What's actually happening — the governance failure

The mechanism to understand at Nidec is not share gain at a key customer; it is the anatomy of the accounting failure and what it implies. The Third-Party Committee's findings describe misconduct that is not a single isolated error but a "wide range of instances" spanning multiple years and multiple types — inventory not written down, asset values falsely recognised — concentrated in but not limited to the automotive business.

The restatements going back to the fiscal year ended March 2022 mean that four-plus years of reported earnings, against which investors made decisions, were wrong. The roughly ¥250 billion impairment is the visible cost; the invisible cost is that the market can no longer take any Nidec number at face value until a clean audit re-establishes a baseline.

The Executive Responsibility Investigation Committee — Nidec investigating whether its own current and former directors and auditors are legally liable — is an extraordinary step, and it cuts both ways. It signals the board is treating the matter seriously, which is necessary for eventual rehabilitation. But it also confirms the failure reached the top of the organisation, and it guarantees a prolonged period of management distraction, potential litigation, and leadership uncertainty. For a company built on an aggressive acquire-and-integrate model under a famously dominant founder-leader culture, an accounting scandal of this reach raises the deeper question of whether the reported success of that model was ever what it appeared to be. That is the question an investor cannot answer from the outside today — which is the entire reason the name is, for now, uninvestable as a core position.

The competitive threat / the credibility gap

Nidec's competitive threats in motors are real — Chinese motor makers, the broader EV-traction-motor competitive set, and the same low-cost humanoid-actuator supply chain contesting every name in this batch — but they are second-order. The first-order "competitive" problem is the credibility gap relative to its own peers. While Nidec spends 2026 restating accounts, forgoing its dividend, and investigating its executives, the cleaner names in the joint layer — MinebeaMitsumi, Nabtesco, Harmonic Drive Systems — are publishing audited results, raising guidance, and competing for humanoid design wins with management teams that are not distracted by a governance crisis. Capital, talent and customer trust all flow to credibility, and Nidec has, for now, lost that ground.

If there is litigation active here, it is internal and self-directed: the Executive Responsibility Investigation Committee is the relevant body, there is no confirmed external court timeline as of mid-May 2026, and the resolution timing is genuinely unknown. The outcome that would *contain* the damage is a fast, clean restatement, a credible new audited baseline, demonstrable governance reform, and the impairment proving to be bounded near the ~¥250 billion figure. The outcome that would *accelerate* the damage is further impairments beyond ¥250 billion, additional restatement scope, regulatory action, or a delisting-adjacent reporting failure. An investor sizing this name is, in effect, taking a position on a corporate-governance event with an unknown resolution date — not on

motor technology.

The terminal risk

Nidec's terminal risk is not a technology transition — electric motors are not going away, and Nidec's engineering is not the problem. The terminal risk *is* the governance failure itself: that the accounting scandal, the multi-year restatements, the executive-liability probe and the ~¥250 billion (and possibly growing) impairment permanently impair the multiple the market is willing to pay and permanently damage management credibility, so that even a fully rehabilitated Nidec trades at a structural governance discount for years. There is also a real tail risk that the rehabilitation does not go cleanly — that further misconduct surfaces, that the impairment balloons, or that the reporting failures escalate toward regulatory or listing consequences. Nidec's credible path back exists — a clean restated audit, governance overhaul, bounded impairment, and a refocused operating company with its genuinely strong motor franchises intact — but that path is long, the timeline is unknown, and underwriting it today is closer to speculation than investment. The humanoid-motor capability is real and could matter eventually; it cannot rescue a thesis whose binding constraint is whether the company's numbers are true.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. The motor capability and manufacturing scale are genuinely best-in-class. Nidec is the world's largest electric-motor maker with deep compact-precision and frameless-motor engineering — exactly what humanoid actuators require — and the scale to make them at cost. The industrial logic for Nidec as a humanoid-motor supplier is sound; the problem is not the product.

2. A clean restatement could be a sharp clearing event. If Nidec delivers a fast, credible restated audit with the impairment bounded near ~¥250 billion and demonstrable governance reform, the uncertainty discount could compress quickly — distressed-recovery situations re-rate hard when the binary resolves favourably.

3. The diversified motor franchises outside automotive are likely sound. The appliance, commercial, industrial and small-precision-motor businesses are not where the misconduct concentrated; a refocused Nidec built around those, with the automotive impairment taken, could be a legitimately investable industrial.

4. The valuation already reflects distress. At a depressed share price with the dividend cut and the scandal public, much bad news is in the price — this is a recovery option, not a momentum chase, for the investor who can tolerate the risk profile.

Gap

1. There is no reliable set of financial statements. Nidec has delayed FY-March-2026 results and is restating back to FY-March-2022. The 14.2x forward P/E is computed off numbers the company itself has disavowed. You cannot underwrite a business whose accounts are under reconstruction.

2. The accounting misconduct is broad and multi-year. The Third-Party Committee found a "wide range of instances" across years — worthless inventory carried at value, falsely recognised asset values. This is a pattern, not an error, and it reached senior levels.

3. The impairment is large and may grow. Roughly ¥250 billion in expected impairments, primarily automotive goodwill and fixed assets, with the company explicitly warning of possible additional impairments as derivative impacts. The downside is not yet bounded.

4. Management is distracted and its credibility is impaired. The Executive Responsibility Investigation Committee — Nidec probing its own directors and auditors for legal liability — guarantees prolonged leadership uncertainty and distraction precisely while cleaner competitors compete for humanoid design wins.

5. The resolution timeline is unknown and slippage is itself bearish. There is no confirmed date for the restated statements; the ~June 30, 2026 placeholder is an estimate. Every further delay compounds the credibility damage and the opportunity cost.

Optionality

Event	Date / window	Direction
Release of restated / delayed FY-March-2026 statements	~Q2-Q3 2026 (date unconfirmed)	Binary — the entire thesis gates here
Final impairment quantification	2026	Bear if it exceeds ~¥250 billion
Executive Responsibility Investigation Committee findings	2026-2027	Bear if it triggers litigation / regulatory action

Event	Date / window	Direction
Governance-reform plan and new audited baseline	2026-2027	Bull if credible and clean
Any regulatory / listing-status development	Ongoing	Bear — tail risk worth monitoring

---

The trade

Nidec is, on engineering merit, a legitimate humanoid-motor name — and it is, on governance reality, uninvestable as anything more than a tiny speculative recovery position right now. That is why it is Bucket D, conviction 3: not because the motor business is bad, but because an investor has no reliable instrument to underwrite it. The disciplined posture is to *not* take a position pending resolution; for the risk-tolerant investor who insists on a toehold, cap it at 0.5% of risk capital — speculation sizing, not investment sizing — within a JPY 2,553-2,821 range (current JPY 2,687 ±5%, noting this price is itself anchored to a scandal-impaired tape), stop at JPY 2,200, and treat the release of the restated financial statements (placeholder ~June 30, 2026, unconfirmed) as the binary that determines whether this becomes investable at all. The cleaner expression of the "Japanese motor and joint-layer into humanoid" thesis is emphatically elsewhere in this batch: MinebeaMitsumi (6479) gives you compact-motor and micro-actuator humanoid exposure at 18x with audited accounts, a raised guidance, and no governance cloud — there is no good reason to take Nidec's credibility risk to get motor exposure when a clean alternative trades cheaply. Revisit Nidec only after the restated audit is in hand. **Conviction: 3 / 10.**

Sources referenced inline throughout: Nidec Third-Party Committee investigation report disclosures (February 27, 2026) and final report (April 17, 2026); Nidec announcements on delayed FY-March-2026 results, the ~¥250 billion expected impairment, the year-end dividend forgone, and the Executive Responsibility Investigation Committee. All pre-restatement financial figures are treated as unreliable per the company's own disclosures. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

Section 5 — Layer 2 — Senses: vision, depth, touch

The senses layer is where the robot understands the world. Three modalities, in ascending order of difficulty: vision (2D and 3D), inertial (proprioception via accelerometers/gyros), and force/tactile (the dexterity bottleneck).

Vision is mature; touch is where the trade is. Cognex (CGNX) and Keyence (6861.T) are the global leaders in 2D industrial machine vision, with thirty-year track records and high-margin oligopoly economics. Three-dimensional depth perception is more fragmented: stereo and structured-light cameras are largely captive (each robot OEM builds its own); time-of-flight lidar leadership belongs to Ouster (OUST) on the Western side; FMCW lidar — which returns velocity per pixel along with distance, useful for autonomous-vehicle perception and industrial inspection — is Aeva's (AEVA) pure-play domain; imaging radar, the all-weather modality, is Arbe Robotics' (ARBE) lead. Vision-processing SoCs (the chips that turn raw pixels into features) split between Ambarella's CVflow architecture and indie Semiconductor's (INDI) multimodal SoC family. Mobileye (MBLY) is the autonomous-vehicle perception leader (EyeQ family is in 800 million cars) and now an explicit humanoid play via its 2026 acquisition of Mentee Robotics.

Touch is the dexterity unlock. Without force and tactile sensing at the wrists, fingers and joints, a humanoid is blind to contact: it does not know when its hand has touched a surface, how hard it is gripping a cup, whether it has dropped a tool. Vishay Precision Group (VPG) has the early lead in humanoid tactile and joint-torque sensing — strain-gauge-based sensors are already shipping to multiple humanoid developers with order books ramping toward hundreds of robots per week by end-2026. Sensata Technologies (ST) is the broader force/torque/position sensor supplier (more auto-exposed than humanoid today). Allegro MicroSystems (ALGM) is the leading magnetic position-sensing chip supplier — its chips are the named encoders inside multiple public humanoid BOMs, and the company has the calmest chart in this layer (RSI 55, +19% vs 50-DMA — unextended relative to the rest of the universe).

What changed post-PhotonCap May 19. Nothing tactical, but the layer's relative

importance moved up in our framework. The mention concentration on VPG, Allegro, Mobileye, Ouster and Aeva in his piece corroborates the layer's emerging status as the second-most-concentrated cost slice (after the joint layer). Our existing reads — CGNX and ALGM as the Tier-1 BUY anchors, VPG as a high-conviction WATCH gated on entry timing (RSI 90, +95% vs 50-DMA — extraordinarily extended despite the thesis being right) — are unchanged.

ALGM — Allegro MicroSystems, Inc. · BUY (Tier-1) · Conv 7/10 · Bucket B

Allegro MicroSystems, Inc. (ALGM)

The #1 magnetic position-sensing chip company — robot joints and motors are a fast-doubling option on top of a recovering auto and industrial base.

Investment Research · Photoncap-style deep dive · v1 of "Allegro MicroSystems" · May 14, 2026

What Allegro MicroSystems physically does

Allegro makes the semiconductor chips that tell a motor where it is and how hard it is working. At the physics level, the core product is the magnetic sensor IC — a chip that detects magnetic fields using Hall-effect or magnetoresistive (TMR/GMR) sensing elements integrated with signal-conditioning circuitry on a single die. Pair that chip with a small magnet on a rotating shaft and you have a position sensor: the chip reads the changing magnetic field as the shaft turns and outputs precise angular position. Pair it with a current-carrying conductor and you have a current sensor: the chip reads the magnetic field the current generates and outputs the current value. Allegro's two reporting lines are Magnetic Sensors and Power Products (gate drivers, regulators, motor drivers) — and the magnetic-sensing franchise is the crown jewel, where Allegro holds roughly 28% global share in magnetic sensor ICs, the largest single position in the segment, per early-2026 market data.

Why this is a binding constraint for robotics: every joint in a robot is a motor, and every motor needs to know its position to be controlled precisely. A humanoid robot has dozens of actuated joints — each one needs angular position sensing for the joint and current sensing for the motor that drives it. This is not optional and it is not a "good enough with a cheap part" function: precise, low-latency, drift-stable position and current sensing is what separates a robot that moves smoothly and safely from one that does not. Allegro's magnetic sensor ICs are designed for exactly the automotive-grade

reliability and precision that humanoid joints require — which is why robotics design wins are landing in its lap rather than going to commodity sensor makers.

The honest framing: Allegro is an automotive semiconductor company first — roughly 71% of revenue is auto — with a fast-growing industrial-and-robotics franchise layered on top. You are not buying a robotics pure-play; you are buying the magnetic-sensing leader at the moment robotics is becoming a visible, doubling growth driver on top of a recovering core business.

Product roadmap

Allegro's roadmap runs along two axes: the magnetic-sensing technology transition and the application expansion into robotics and data center. On the technology side, the key move is the shift from Hall-effect to TMR (tunneling magnetoresistance) sensing — TMR offers higher sensitivity, better signal-to-noise and lower power, which matters for the precision and efficiency demands of robot joints and high-speed power electronics. The marquee recent product is the ACS37100, launched October 2025 — Allegro describes it as the first commercial 10 MHz TMR magnetic current sensor, designed to pair with high-speed wide-bandgap (silicon-carbide and gallium-nitride) power electronics, with a 50-nanosecond response time and low noise. This is the part that positions Allegro for the next generation of efficient motor drives, which is exactly the technology stack a modern robot actuator uses.

On the application side, the roadmap is the deliberate push beyond automotive. Allegro's "Industrial and Other" segment now explicitly targets data center (power and current sensing for AI server racks) and robotics (joint position and motor current sensing). Management disclosed on the Q4 FY2026 call (May 7, 2026) that it secured two design wins with leading Chinese robotic companies for robotic joints during the quarter, and that robotics-and-automation sales doubled year-over-year in fiscal 2026. The product approach is to take automotive-qualified magnetic-sensing IP — already proven to automotive reliability standards — and configure it for robotic joints, which is a faster path to a credible product than building from scratch. What Allegro does not make is the motor, the actuator assembly, the encoder module, or the robot — it makes the sensing silicon inside those systems. It also is not a force/torque or tactile-sensing company; that is VPG's and Sensata's territory. Allegro's lane is position and current.

The financial print

Allegro closed fiscal 2026 (year ended late March/early April 2026) with full-year sales of $890 million, up 23% year-over-year, and non-GAAP EPS of $0.54, more than double the prior year — results released May 7, 2026. The Q4 FY2026 print was the fifth consecutive quarter of sales growth at $243 million, with non-GAAP EPS of $0.17, nearly triple year-over-year. The segment detail is where the robotics story lives: Q4 Magnetic Sensor revenue was $141 million (up 21% YoY), Power Products revenue was $102 million (up 35% YoY), and Industrial and Other sales were $79 million, up 49% YoY and explicitly "led by data center and robotics." For the full year, Industrial and Other sales were $261 million, up 38%, with data center revenue more than quadrupling to about 10% of total sales. Full-year automotive sales were $629 million, up 17%, at 71% of total revenue.

The recovery story is real: Allegro spent 2023–2024 working through an auto/industrial inventory correction, and FY2026's 23% revenue growth and doubled EPS mark the cyclical recovery and the structural data-center/robotics layering landing together. Forward consensus for FY2027 has revenue continuing to grow at a solid double-digit pace with non-GAAP EPS expanding further as operating leverage returns — sell-side coverage (the names active on ALGM include Morgan Stanley, BofA, Deutsche Bank and others) clusters around continued double-digit top-line growth and meaningful margin recovery. At $45.44 the stock trades at a forward P/E of 31.0 with a market cap of $8.4B — the cheapest forward multiple of the four growth names in this batch (VPG 66x, CGNX 37x, Keyence 41x), on a company with a genuine robotics design-win story. The next binary is Q1 FY2027 (the June 2026 quarter) earnings, expected around August 6, 2026.

Customer mix today

Allegro's revenue mix is disclosed cleanly by end market. Automotive was roughly 71% of fiscal 2026 revenue at $629 million, up 17% — within that, "Focused Auto" sales (Allegro's term for its targeted high-content applications) were $349 million, up 30%, accounting for 55% of total auto. Industrial and Other was roughly 29% of revenue at $261 million, up 38% — and inside that bucket, data center more than quadrupled to about 10% of total company sales, while robotics-and-automation sales doubled year-over-year (Q4 FY2026 call, May 7, 2026).

The 2024-versus-2026 change is the story. Two years ago Allegro was a roughly four-fifths-automotive company in the depths of an inventory correction, with industrial as a sleepy minority. By FY2026 the Industrial and Other segment is the fastest-growing line, data center has gone from negligible to ~10% of total sales, and robotics has emerged as a doubling, design-win-driven contributor. The customer base is not disclosed by name

— Allegro sells to tier-1 automotive suppliers, OEMs, industrial-equipment makers, and now Chinese robotics companies — but the structural shift is clear: the growth is moving from cyclical auto recovery toward secular industrial, data-center and robotics demand. Robotics is still a small slice of the 29% Industrial-and-Other bucket — best read as low-single-digit percent of total revenue today — but it is doubling annually and the design wins are landing now, which is precisely the early-cycle profile this theme is looking for.

What’s actually happening in robotics design wins

The mechanism tying Allegro to the robotics theme is the joint count. A humanoid robot has dozens of actuated joints; each joint needs magnetic position sensing for angle and magnetic current sensing for the motor. Allegro’s path into this socket is its automotive-qualified magnetic-sensing IP — parts already proven to the reliability, precision and temperature-stability standards that automotive demands, which happen to be the same standards humanoid joints need. That gives Allegro a credible-product head start over commodity sensor makers and a faster qualification path than building robotics-specific silicon from scratch.

The concrete proof points from the Q4 FY2026 call: Allegro secured two design wins with leading Chinese robotic companies for use in robotic joints during the quarter, and robotics-and-automation sales doubled year-over-year in fiscal 2026, with management citing “growing engagement and design wins with humanoid robotics customers.” The China angle is notable — China is where humanoid robot volume is ramping fastest and where the OEM ecosystem is most willing to design in merchant silicon rather than vertically integrate, which favors a merchant sensor supplier like Allegro. The skeptical read: two design wins and a doubling-off-a-small-base is real but early, and design wins are not revenue until the robots ship in volume. The constructive read: this is exactly what the early innings of a structural socket-capture looks like — qualified product, design wins landing, the customer base expanding from auto into a new high-joint-count application — and Allegro is capturing it at automotive-grade margins rather than commodity ones. The number to watch is whether the “robotics doubled” line keeps doubling and whether the design-win count keeps climbing on the FY2027 calls.

The competitive threat / Infineon and Melexis

Allegro’s named direct competitors in magnetic position and current sensing are Infineon, Melexis, TDK, Asahi Kasei (AKM) and Analog Devices. Allegro’s roughly 28%

global share in magnetic sensor ICs is the largest single position, but this is a genuine oligopoly, not a monopoly — Infineon in particular is a formidable competitor that, in January 2025, consolidated its sensor and RF operations into a dedicated Sensor Units & Radio Frequency (SURF) business unit, signaling strategic focus on exactly this market. Melexis is the other pure-play threat, deeply embedded in automotive position sensing. TDK brings materials and miniaturization strength.

There is no active IP litigation to flag. The competitive dynamic in the robotics socket specifically is a race to qualify: every magnetic-sensor vendor sees the humanoid joint-count opportunity, and the question is who lands the design wins as the humanoid OEMs lock their bills of materials. Allegro's advantages are its share leadership, its TMR roadmap (the ACS37100 and the broader TMR transition give it a technology edge on the highest-precision sockets), and its automotive-qualification credibility. The risks are that Infineon's scale and SURF focus let it out-invest Allegro, and that the Chinese robotics OEMs — where Allegro is winning today — may over time favor domestic Chinese sensor suppliers as that ecosystem matures, a geopolitical-localization risk that is hard to quantify but real. The bear case via competitor is not that Allegro loses its leadership outright, but that the robotics socket gets split four or five ways and Allegro captures a smaller share of it than its 28% magnetic-IC position would suggest.

The terminal risk

The terminal risk for Allegro is an architectural shift away from magnetic sensing in robot joints. Magnetic position sensing is the incumbent approach, but it is not the only one — optical encoders, inductive (eddy-current) encoders and capacitive encoders all have credible proponents for joint position sensing, and some humanoid designs use them. If the humanoid industry standardizes on a non-magnetic encoder architecture for joints — for cost, precision, or integration reasons — Allegro's automotive-magnetic-sensing incumbency does not transfer, and the robotics growth leg the bulls are paying for thins out. The second terminal path is plainer commoditization: magnetic sensor ICs are a maturing semiconductor category, and if the robotics socket turns into a price war among five qualified vendors, the automotive-grade margin Allegro expects to earn there compresses toward commodity levels.

Neither is imminent, and Allegro has structural defenses. Its core business is ~71% automotive — a market where magnetic sensing is deeply entrenched and not going anywhere — so the terminal risk applies to the robotics growth option, not the base business. Its TMR roadmap keeps it on the precision frontier where commoditization is slowest. And the breadth of its end markets (auto, data center, industrial, robotics) means no single architectural shift breaks the company. The constraint on the multiple:

at 31x forward earnings, Allegro is the cheapest growth name in the batch precisely because the market is not yet fully pricing the robotics option — which means the terminal risk is more “the robotics upside doesn’t materialize” than “the company is impaired.” That asymmetry is part of what makes the name attractive.

Bull / Gap / Optionality (Photoncap framing)

1. Cheapest forward multiple in the batch on a real robotics story. At 31.0x forward earnings, Allegro trades well below VPG (66x), Keyence (41x) and Cognex (37x) — yet it has concrete robotics design wins (two with leading Chinese robotic companies for robotic joints, Q4 FY2026 call) and a robotics-and-automation line that doubled year-over-year in FY2026. The robotics option is partly unpriced.

2. Market-share leadership in the exact sensing category robots need. Allegro holds roughly 28% global share in magnetic sensor ICs, the largest single position (early-2026 market data). Every robot joint needs magnetic position and current sensing — Allegro is the incumbent leader in precisely that socket, with automotive-grade qualification credibility competitors must match.

3. The full-year FY2026 print confirms a real cyclical-plus-structural inflection. FY2026 revenue of $890 million was up 23%, non-GAAP EPS of $0.54 more than doubled, and it was the fifth consecutive quarter of sequential sales growth (company release, May 7, 2026). The auto/industrial inventory correction is over and the data-center/robotics layering is landing on top of the recovery.

4. Data center is the proof-of-concept for the merchant-silicon expansion playbook. Data-center revenue more than quadrupled in FY2026 to ~10% of total sales — Allegro has already demonstrated it can take its sensing/power IP into a new high-growth application and scale it fast. Robotics is the same playbook one cycle behind, which de-risks the robotics ramp.

5. The TMR technology transition is a durable edge. The ACS37100 (launched October 2025, the first commercial 10 MHz TMR current sensor, 50ns response) puts Allegro on the precision frontier for high-speed wide-bandgap motor drives — exactly the actuator technology modern robots use. TMR keeps Allegro ahead on the highest-value, slowest-to-commoditize sockets.

Gap

1. It is ~71% automotive and the robotics line is still small. Automotive was 71% of FY2026 revenue; robotics is a low-single-digit slice of the 29% Industrial-and-Other

bucket. The robotics story is real and doubling, but the stock still trades on the auto cycle in the near term — an auto downturn would dominate the robotics tailwind.

2. Design wins are not revenue. Two robotic-joint design wins and a "doubled YoY" line are an early-innings signal, not a proven revenue stream. Humanoid volume timing is uncertain across the whole theme, and if the production ramps slip, the robotics contribution stays small longer than the multiple assumes.

3. China concentration in the robotics wins cuts both ways. Allegro's robotics design wins are with leading Chinese robotic companies — the fastest-ramping ecosystem, but also the one most exposed to eventual domestic-supplier localization and to US-China trade friction. The robotics growth leg has a geopolitical overhang that is hard to quantify.

4. Magnetic sensing faces both architectural and pricing competition. Optical, inductive and capacitive encoders are credible non-magnetic alternatives for joint sensing, and Infineon (with its 2025 SURF business-unit focus), Melexis, TDK and ADI are all chasing the same robotics socket. Allegro's leadership is not guaranteed to translate into a leading share of the robotics socket specifically.

Optionality

Event	Date / window	Direction
Q1 FY2027 (June quarter) earnings	~Aug 6, 2026	Binary on the recovery + robotics cadence
Robotics-and-automation line keeps doubling / design-win count climbs	FY2027 calls	Bull — confirms the socket capture
Additional humanoid OEM design wins (beyond the two Chinese names)	H2 2026	Bull if delivers
Data-center revenue sustains its quadrupling trajectory	FY2027	Bull — validates the expansion playbook
Automotive end-market cyclical wobble	Ongoing	Bear — dominates near-term given 71% mix

---

The trade

Allegro is the best risk/reward in the sensing batch — a genuine robotics design-win story available at the cheapest forward multiple of the four growth names, on the least-extended chart, with a recovering core business underneath it as a floor. The structural case: the #1 magnetic position-sensing chip company, capturing the exact joint-sensing socket every robot needs, with two robotic-joint design wins already booked, a robotics line doubling year-over-year, and the data-center expansion as proof the merchant-silicon playbook works. The chart is the calmest in the batch — RSI 55.2 and +19.3% above the 50-day average is genuinely investable, not a chase. The trade is to initiate in a zone of current ±5%, roughly $43.17–$47.71, accumulating on any dip toward the lower end; size at 1.0–1.5% of risk capital — a real position, justified by the valuation cushion, the recovering base business, and the unpriced robotics option. Stop at roughly $37.50, below the structural base and the rising 50-day cloud; a break there would say the auto/industrial recovery has stalled. The named catalyst is Q1 FY2027 (the June 2026 quarter) earnings around August 6, 2026, where the post-recovery growth cadence and the robotics design-win progression get their first test of the new fiscal year. If you want a higher-torque but far more extended expression of the same robotics-sensing theme, VPG (this batch) is the pure-play — but it trades at double the multiple on a 90+ RSI; Allegro is the disciplined way to own the sensing layer. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

CGNX — Cognex Corporation · BUY (Tier-1) · Conv 7/10 · Bucket B

Cognex Corporation (CGNX)

The undisputed leader in industrial machine vision — "the robot's eyes" — seven quarters into a margin turnaround and freshly re-platformed on edge AI.

Investment Research · Photoncap-style deep dive · v1 of "Cognex" · May 14, 2026

What Cognex physically does

Cognex makes the systems that let a machine see and make a decision about what it sees. A machine-vision system is, at the hardware level, an industrial camera with a lens, controlled lighting, an image sensor and an onboard processor — but the value is in the software stack that runs on it. Cognex's software takes the raw image and answers a question a factory or warehouse cares about: is this part defective, what code is on this box, where exactly is this object so a robot can pick it. The company's product DNA is two software families — VisionPro (the PC-based vision software platform) and In-Sight (self-contained "smart cameras" that run vision software on the camera itself) — plus the ID products (barcode and code-reading, the DataMan family) and the deep-learning suite (historically ViDi, now folded into the broader AI vision offering).

Why this is the binding constraint for robotics: a robot cannot manipulate what it cannot locate. Vision-guided robotics — where a camera tells the robot arm where the object is, in three dimensions, in real time — is the function that turns a blind, pre-programmed arm into something that can handle variation. Cognex sits at exactly that chokepoint. Its machine-vision systems are the perception layer between the camera and the robot controller, and in a world moving toward general-purpose robots that operate in unstructured environments, the vision problem gets harder and more valuable, not easier.

The structural shift inside Cognex's own technology is the move to edge AI. Traditional machine vision was rule-based — an engineer programmed the system to look for specific features. The new generation runs neural-network inference on the camera itself, so the system can be trained on examples rather than hand-programmed, and can handle the messy variation that rule-based vision could not. Cognex's Q1 2026 product

launches — the In-Sight 6900 powered by NVIDIA and the In-Sight 3900 powered by Qualcomm — are the company re-platforming its entire smart-camera line onto AI silicon. This is the bet that the next decade of factory and warehouse vision is AI-native, and that Cognex's installed base and software moat carry it across that transition.

Product roadmap

Cognex's roadmap in 2025–2026 is dominated by the edge-AI re-platforming. The In-Sight smart-camera family — historically the In-Sight 8000 series (2D smart cameras) and the higher-end vision systems — is being rebuilt around AI accelerators. The In-Sight 8900 brought embedded AI to OEM deployments. The flagship 2026 launches came in early May: the In-Sight 6900, an AI vision system powered by NVIDIA silicon, and the In-Sight 3900, an embedded AI vision system launched May 5, 2026 on Qualcomm's Dragonwing platform — Cognex says the 3900 delivers up to 4x faster processing than prior Cognex generations, supports up to 25-megapixel imaging, and runs PC-free for real-time edge inspection. On the ID/code-reading side, the DataMan 290 (recent) added AI-enabled auto-setup and advanced code filtering.

The cadence shows a company pushing AI down its entire price stack — from high-end NVIDIA-powered systems for demanding applications to Qualcomm-powered embedded units for cost-sensitive OEM volume. At its June 2025 Investor Day and a March 2026 conference appearance, management laid out the strategic frame: AI-driven growth, deliberate exits from low-margin legacy product lines, and a raised EBITDA-margin target of 25–31%. What Cognex does not make is the robot, the actuator, or the end-to-end robot software — it is the perception layer, sold to robot integrators, OEMs and end users. It also is not a 3D-vision-only or a camera-component-only player; its differentiation is the full hardware-plus-software-plus-AI stack, sold across 30-plus industries.

The financial print

Cognex closed FY2025 (year ended December 31, 2025) with revenue of $994 million, up 9% year-over-year (up 8% constant-currency, or 7% excluding a one-time medical-lab channel-partner benefit), and a full-year operating margin of 16.3% — results released February 11, 2026. Adjusted EBITDA margin was 21.5%, up 440 basis points year-over-year, clearing the company's 20% milestone ahead of plan. The story is the margin trajectory: Cognex spent 2022–2024 over-earning then under-earning through the post-COVID factory-automation cycle, and FY2025 was the year the cost discipline and mix

improvement showed up.

The Q1 2026 print on May 6, 2026 accelerated it. Revenue was $268.4 million, up 24.3% year-over-year and ahead of expectations, with double-digit growth across logistics, packaging, electronics and semiconductor. Operating margin was 22.3% and adjusted EBITDA margin reached 26.9%, up 1,010 basis points year-over-year — the seventh consecutive quarter of margin improvement. Adjusted diluted EPS was $0.34, up 113% year-over-year, the seventh straight quarter of EPS growth. Management guided Q2 2026 to revenue of $280–300 million (midpoint ~16.5% growth) and adjusted EPS of $0.40–0.44 (midpoint implying ~68% growth). At $64.61 the stock trades at a forward P/E of 37.2 with a market cap of $10.75B — a premium multiple, but on a company that is compounding both revenue and margin, which is a different proposition than a flat-revenue name on a high multiple. The next binary is Q2 2026 earnings, expected around July 30, 2026.

Customer mix today

Cognex breaks its revenue out by end market, and the FY2025 mix is well-disclosed. Logistics was roughly 26% of full-year 2025 revenue and delivered its ninth consecutive quarter of double-digit growth — this is the warehouse-automation engine, driven by e-commerce parcel sortation and the structural build-out of automated distribution. Packaging was roughly 21% of 2025 revenue and posted double-digit Q1 2026 growth, strong enough that management raised its 2026 packaging outlook to high-single-digit. Consumer electronics was roughly 19% of revenue, with double-digit Q1 growth and a high-single-digit to double-digit full-year 2026 outlook — this is the historically lumpy Apple-supply-chain exposure that whipsaws on smartphone build cycles. Automotive sat in the low teens of revenue and was the one soft spot, down high single digits in 2025, though management said on the Q4 2025 call that the market has bottomed and guided flat-to-low-single-digit growth for 2026.

The 2024-versus-2026 change worth highlighting is the diversification away from automotive and consumer-electronics lumpiness toward the more durable logistics and packaging demand, plus the layering-on of the AI-vision product cycle. Cognex does not have the single-customer concentration of a supplier — its largest customers are integrators and end users spread across 30-plus industries — which is a structural strength versus several names in this batch. The robotics/humanoid-specific exposure is not yet broken out as its own line; it is embedded across logistics (warehouse robots) and the broader vision-guided-robotics demand, and it is best understood as an option on the theme rather than a quantified revenue stream today.

What’s actually happening in logistics and AI vision

The two mechanisms driving Cognex right now are the logistics flywheel and the AI-vision product transition. On logistics: nine consecutive quarters of double-digit growth is not a cycle, it is a structural build — every automated distribution center needs barcode reading at every sortation point and increasingly needs vision-guided picking and induction. Cognex’s installed base in logistics compounds because once a DC standardizes on Cognex code-reading, the expansion revenue is sticky. Management was deliberately conservative on the 2026 logistics outlook (mid-to-high single digit) despite the Q1 strength, which reads as sandbagging given the run-rate.

On AI vision: the In-Sight 6900 (NVIDIA) and In-Sight 3900 (Qualcomm, launched May 5, 2026) are the proof points of the re-platforming. The strategic logic is that AI-native vision expands Cognex’s addressable market — applications that were impossible or uneconomic with rule-based vision (handling high variation, training on examples rather than code) become solvable, and Cognex can sell into them. The seven consecutive quarters of margin expansion suggest the transition is being executed without the margin destruction that usually accompanies a platform shift, because Cognex is simultaneously exiting low-margin legacy lines (the “portfolio exits” management flagged in March 2026). The skeptical read: AI vision is also the door through which competition and commoditization walk in, which is the terminal-risk section. The bull read: Cognex’s software moat and installed base are exactly what carry an incumbent across a platform transition, and the early evidence — 24% revenue growth, 27% EBITDA margin in Q1 — says it is working so far.

The competitive threat / Keyence

Cognex’s primary named competitor is Keyence (6861, also covered in this batch). Per coherent market-share estimates, Cognex holds roughly 21% of the global machine-vision market and Keyence roughly 19% — together close to half the market, with the remaining leading-five players controlling another slice and a long tail of Teledyne (DALSA), Basler, Sick, Omron, Datalogic, Canon and others. Keyence is the more formidable competitor: it is structurally more profitable (operating margins above 50% versus Cognex’s mid-teens-to-low-20s), it dominates the Asia-Pacific market, and its direct-sales model gives it deep application-engineering reach. Where Cognex is differentiated is depth of vision software and the AI/deep-learning stack — VisionPro

and the AI suite are genuinely best-in-class — and a stronger position in logistics and code-reading specifically.

There is no active IP litigation to flag between the two. The competitive dynamic is a share contest, not a courtroom one: both companies are racing to AI-native vision, and the question is whether Cognex's software lead or Keyence's profitability and sales reach wins the next product cycle. The honest assessment is that this is a structural duopoly-plus where both can grow with the market — machine vision is a ~$12B market in 2026 growing mid-to-high single digits — but Keyence's margin structure means it can out-invest and out-price Cognex if it chooses to. Cognex's defense is that it is the vision-software specialist while Keyence is a broad-line sensor company for which vision is one of many products; in a robotics world where vision software is the hard part, that specialization should matter.

The terminal risk

The terminal risk for Cognex is that general-purpose AI eats dedicated machine vision. The transition technology is the vision-language model — large multimodal AI models that can look at an image from a commodity camera and answer perception questions without Cognex's specialized hardware or software. If a warehouse can run a VLM on a cheap camera and a standard GPU and get "good enough" defect detection, code reading and object localization, the value of a dedicated Cognex system erodes. This is the SaaS-displaced-by-LLM analogy applied to industrial vision: the moat was the specialized software, and general-purpose AI is a solvent for specialized software.

The timing is genuinely uncertain. The skeptic case says this is a 3–5 year threat as VLMs get faster, cheaper and reliable enough for industrial uptime requirements. The bull rebuttal — and it is a real one — is that industrial vision has requirements that general-purpose AI does not naturally meet: deterministic latency, 99.9%+ reliability, traceability for regulated industries, and integration with PLCs and robot controllers. Cognex's response is to be the company that productizes AI for industrial requirements rather than the company that gets displaced by it — the In-Sight 6900/3900 launches are precisely that bet, putting NVIDIA and Qualcomm AI silicon inside ruggedized, deterministic, factory-grade hardware. Named alternative beneficiaries of the displacement scenario would be NVIDIA's own robotics/vision stack, hyperscaler vision APIs, and AI-native vision startups. The constraint this puts on the multiple: 37x forward earnings already prices Cognex as a durable AI winner, so the terminal risk is less "the stock goes to zero" and more "the multiple compresses if the AI-native transition turns out to commoditize the category."

Bull / Gap / Optionality (Photoncap framing)

1. Seven straight quarters of margin expansion, and accelerating. Q1 2026 adjusted EBITDA margin of 26.9% was up 1,010 basis points year-over-year (company release, May 6, 2026), and management has raised the long-term EBITDA target to 25–31%. This is a structural re-rating of the earnings power, not a cyclical bounce — the cost discipline and portfolio exits flagged in March 2026 are durable.

2. The logistics flywheel is a structural, not cyclical, demand engine. Nine consecutive quarters of double-digit logistics growth (Q1 2026 call), with management still guiding conservatively to mid-to-high single digit for 2026 — the sandbag is visible. Logistics at ~26% of revenue is the most durable leg of the mix and compounds on installed-base expansion.

3. The AI-vision re-platforming expands the addressable market. The In-Sight 6900 (NVIDIA) and In-Sight 3900 (Qualcomm, launched May 5, 2026) make previously-uneconomic applications solvable. Cognex is putting AI silicon across its full price stack while exiting low-margin legacy lines — the cleanest version of an incumbent navigating a platform shift.

4. Best balance-sheet and customer-concentration profile in the batch. Cognex sells across 30-plus industries with no single-supplier-style customer concentration, carries a strong net-cash position, and is the global market-share leader. In a batch that includes a thin-margin turnaround (VPG) and an auto-heavy cyclical (Sensata), Cognex is the quality name.

5. Robotics is free optionality on top of a working base case. The base case — 24% Q1 revenue growth, margin expansion, logistics durability — does not require the humanoid theme to work. Vision-guided robotics and humanoid perception are an unpriced call option layered on a company that is already compounding.

Gap

1. The multiple already prices a durable AI winner. At 37.2x forward earnings, Cognex is not cheap on any absolute measure — the market is paying for the margin trajectory and the AI transition both succeeding. RSI 68.2 and +19.8% above the 50-day average is extended (though far less so than VPG); a growth disappointment re-rates this meaningfully.

2. Consumer electronics and automotive are still lumpy. Roughly 19% of revenue is consumer electronics tied to smartphone build cycles, and automotive (low-teens of

revenue) was down high single digits in 2025. Both are guided to recover in 2026, but a slip in either would expose how much of the Q1 strength was end-market timing.

3. The VLM commoditization threat is real and Cognex is partly betting against itself. By moving to AI-native vision, Cognex is making its own category more accessible to general-purpose AI competitors. The In-Sight 6900/3900 strategy is the right one, but it is a bet that industrial requirements protect the moat — and that bet is not yet proven over a full cycle.

4. Keyence is a structurally stronger competitor on margin and reach. Keyence runs 50%+ operating margins versus Cognex's low-20s and has the financial capacity to out-invest or out-price Cognex in any segment it prioritizes. Cognex's software lead is real but it competes against a peer that is simply more profitable per dollar of revenue.

Optionality

Event	Date / window	Direction
Q2 2026 earnings vs. $280–300M revenue / $0.40–0.44 EPS guide	~Jul 30, 2026	Binary on margin trajectory
In-Sight 6900 / 3900 AI-vision adoption traction	H2 2026	Bull if revenue contribution shows
Consumer electronics build-cycle recovery	H2 2026	Bull if smartphone demand firms
Automotive end-market inflects from down-HSD to growth	FY2026	Bull — confirms the bottom call
A credible VLM-based industrial-vision deployment at scale	2026–2028	Bear — terminal-risk signal

---

The trade

Cognex is the highest-quality name in this sensing/vision batch and the one whose base case does not depend on the humanoid theme working at all. The structural thesis is straightforward: the dominant machine-vision company, seven quarters into a real margin turnaround, re-platforming onto edge AI in a way that expands its market,

with the best balance sheet and the least customer concentration in the batch — and robotics/humanoid perception as free optionality on top. The chart is moderately extended (RSI 68.2, +19.8% versus the 50-day) but nowhere near the VPG-style danger zone, so this is a name you can actually initiate rather than only watch. The trade is to initiate in a zone of current ±5%, roughly $61.38–$67.84, accumulating on any pullback toward the lower end; size at 1.0–1.5% of risk capital — a real position, justified by the quality and the working base case, but tempered by the 37x forward multiple that leaves no room for a growth stumble. Stop at roughly $54, below the structural base and the rising 50-day cloud; a break there would say the margin-expansion narrative has cracked. The named catalyst is Q2 2026 earnings around July 30, 2026, where the $280–300M revenue and $0.40–0.44 EPS guide gets tested and the AI-vision product traction starts to show. There is no cleaner expression of the machine-vision thesis — Cognex is the cleaner expression; the only purer-vision alternative, Keyence, comes with a much higher absolute multiple and Japan-listing friction. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

6861 — Keyence Corporation / · WATCH (Tier-2) · Conv 6/10 · Bucket B

Keyence Corporation (6861.T)

The most profitable sensor and machine-vision company in the world — vision-guided robotics is one growth leg of many, and the 50%+ operating margin is the real moat.

Investment Research · Photoncap-style deep dive · v1 of "Keyence" · May 14, 2026

What Keyence physically does

Keyence makes the sensors, vision systems, measuring instruments and laser markers that factories use to automate inspection and control — and it makes them at a profitability that has no real analogue in the industrial world. At the product level, Keyence's catalogue spans photoelectric and proximity sensors (does a part exist, where is it), measurement sensors (laser displacement, dimensional gauging), machine-vision systems (the CV-X and IV/VS families that inspect and guide), barcode readers, 3D measurement systems, and laser markers. None of this is exotic by itself — competitors make all of these categories. What is exotic is the business model: Keyence is fabless and outsources most manufacturing, runs a direct-sales force of application engineers rather than distributors, and designs products to solve specific customer problems at a price that captures most of the value created. The result is an operating margin above 50% — roughly triple a typical industrial-automation peer and well above even Cognex's low-20s.

Why this matters for robotics: Keyence sits on the perception side of the robot, the same chokepoint as Cognex. Its vision-guided robotics offering — the CV-X series 3D vision-guided robotics, built around a multi-camera-plus-projector architecture, and the VS series that combines high-speed inspection and robot guidance on one platform — is the layer that tells a robot arm where an object is in three dimensions so it can pick, place or assemble it. As factory automation shifts toward higher product mix, higher speeds and tighter tolerances, vision-guided robotics becomes the way factories cope, and Keyence's sensor-and-vision breadth lets it sell the whole perception stack rather

than one component.

The honest framing for a robotics theme: Keyence is not a robotics pure-play and does not pretend to be. It is a diversified factory-automation sensing company for which vision-guided robotics is one important and growing application among many. You are not buying a humanoid-robot bet; you are buying the single best-run sensing franchise in the world, with robotics as a structural tailwind embedded in the broader automation story.

Product roadmap

Keyence's roadmap is best understood as continuous, model-by-model iteration across a very broad catalogue rather than a few flagship launches — the company refreshes its sensor, vision and measurement lines constantly. On the robotics-relevant side, the recent and ongoing work centers on the VS series — a vision system with integrated AI for vision-guided robotics, which combines high-speed inspection and robot guidance on a single platform and cuts setup time — and the CV-X series 3D vision-guided robotics, which uses a four-camera-plus-single-projector design for 3D robot guidance. In 2025 Keyence launched an AI sensor series with built-in machine-learning inference, targeting among other things FDA 21 CFR Part 11-compliant pharmaceutical tablet inspection.

Keyence's patent portfolio (over 500 machine-vision patents per PatSnap's 2026 review) points the development roadmap toward hyperspectral imaging integration, collaborative-robot vision, predictive-maintenance analytics (fusing vision with vibration and thermal data to flag equipment failure 15–30 days ahead), and longer-dated bets on quantum-dot imaging sensors and neuromorphic vision processing for sub-millisecond defect detection. Its adaptive-HDR imaging technology (US patent 12567188B2, issued 2026) is cited at 8–12 ms processing latency per frame. These should be read as a credible roadmap direction, not as dated, confirmed product launches — Keyence is famously tight-lipped and product cadence is rolling, so the roadmap is a trajectory rather than a calendar.

What Keyence does not make is the robot, the actuator, or end-to-end robot software — like Cognex, it is the perception and sensing layer. It also is not a pure machine-vision company; vision is one of roughly half-a-dozen major product categories, which is both a diversification strength and the reason it is not a robotics pure-play.

The financial print

Keyence closed fiscal 2025 (year ended March 2026) with net sales of ¥1.16 trillion (~$7.7B USD at ~150 JPY/USD), a record for the fifth consecutive year and up roughly 10% year-over-year — results released April 27, 2026. Operating income rose about 8% year-over-year, and the operating margin landed in the low-50s percent range — management cited a March-quarter operating margin of 54.3% excluding Cadenas amortization, with the full-year margin near 52%. In the March quarter all overseas markets grew over 20% year-over-year, an acceleration into the fiscal year-end. (Note: some data aggregators garble Keyence's revenue by an order of magnitude — the audited figure is ~¥1.16 trillion, not ¥11.7 trillion.)

The operating-margin number is the entire investment case in one statistic. A ~52% operating margin on a ¥1.16-trillion-revenue industrial company is, simply, the best in the sector globally — for comparison, Cognex runs mid-teens-to-low-20s and a typical automation peer runs high-teens. It reflects the fabless model, the direct-sales value capture, and pricing power that has survived every cycle. Forward consensus has FY2026 (year ending March 2027) net sales continuing to grow at a high-single-digit-to-low-double-digit rate, with the margin holding near current levels — Japanese broker coverage (Nomura, Daiwa, SMBC Nikko, Mizuho among the active houses) clusters around continued double-digit earnings growth, though Keyence does not give detailed guidance. At JPY 77,310 the stock trades at a forward P/E of 41.4 with a market cap of ¥18.7 trillion — a premium, but a premium the market has paid for two decades because the margin and the return on capital justify it. The next binary is Q1 FY2026 (the June 2026 quarter) earnings, expected in early August 2026.

Customer mix today

Keyence's defining customer characteristic is that it has almost no concentration. The company does not disclose — and does not have — a dominant customer or a dominant end market. Its revenue is spread across electronics and semiconductor manufacturing, automotive, machinery, pharmaceuticals and medical, food and beverage, logistics and packaging, and general industrial, sold to tens of thousands of factories across Japan, the Americas, Europe, China and the rest of Asia. No single end market is more than a modest share of the total, and the geographic spread is similarly diversified, with overseas now the larger portion of sales and growing faster than Japan — the March-2026 quarter saw every overseas region up 20%+.

This is a genuine structural advantage and the cleanest customer profile in the entire sensing/vision batch. Where VPG depends on a handful of unnamed humanoid developers and Sensata is roughly 70% automotive, Keyence's revenue base is so diversified

that no single customer loss, no single end-market downturn, and no single geography can break the thesis. The 2024-versus-2026 change worth noting is the continued mix shift toward overseas and toward AI-enabled products — but the headline is the absence of a concentration story, which for a quality-compounder thesis is exactly what you want. The robotics/vision-guided-robotics exposure is embedded inside the machine-vision and sensor categories rather than broken out; it is a growth contributor, not a reported segment.

What’s actually happening in vision-guided robotics

The mechanism that ties Keyence to the robotics theme is the secular shift in how factories handle variation. The traditional factory ran high-volume, low-mix lines with blind, hard-programmed robots. The modern factory runs higher product mix, faster changeovers and tighter tolerances — and the way it copes is vision-guided robotics, where a camera tells the robot where things are in real time so the line can handle variation without re-tooling. Keyence’s VS series and CV-X 3D vision-guided robotics are built for exactly that transition: integrated AI, single-platform inspection-plus-guidance, fast setup.

The specifics that matter: Keyence’s edge here is not a single killer product but the combination of breadth and direct sales. Its application engineers walk into a factory, identify the bottleneck, and can sell whatever combination of sensor, vision system and 3D guidance solves it — and they capture the value of that solution in the price. As the March-2026 quarter’s 20%+ overseas growth shows, the demand is broad and accelerating, not concentrated in one robotics customer or one geography. The skeptical read is that vision-guided robotics is a competitive field and Keyence is not uniquely positioned in any single sub-segment — Cognex is arguably ahead on pure vision software, and the Chinese automation vendors are cheaper. The constructive read is that Keyence does not need to win any single segment; it needs the overall automation-and-robotics adoption curve to keep rising, and to keep capturing its historical share of the value at its historical margin. The evidence — five straight record years, 52% operating margin held through cycles — says it has.

The competitive threat / Cognex

Keyence’s most direct named competitor in the machine-vision portion of its business is Cognex (CGNX, also in this batch). Per coherent market-share estimates, Cognex holds

roughly 21% of the global machine-vision market and Keyence roughly 19% — a structural duopoly-plus at the top, with Teledyne (DALSA), Basler, Sick, Omron, Datalogic and others in the tail. The competitive picture inverts depending on the metric: Cognex leads on depth of vision software and AI/deep-learning tooling and on the logistics/code-reading vertical; Keyence leads decisively on profitability (52% operating margin versus Cognex's low-20s), on Asia-Pacific reach, and on the breadth of the surrounding sensor catalogue that lets it sell the whole perception stack.

There is no IP litigation to flag between the two. Beyond Cognex, Keyence competes across its full catalogue with Omron, Sick, Panasonic, Banner and a long list of category specialists in sensors, and increasingly with lower-cost Chinese automation vendors who attack the price-sensitive end of the market. The competitive risk that actually matters for the thesis is not losing share to Cognex — both can grow with a mid-to-high-single-digit machine-vision market — but the slow erosion of pricing power. Keyence's 52% margin exists because customers pay for the application-engineering value and the time-to-solution. If Chinese vendors close the capability gap, or if general-purpose AI vision makes "good enough" cheap, the margin that underwrites the entire valuation comes under pressure. That is a multi-year risk, not a 2026 one — but it is the bear case worth respecting.

The terminal risk

The terminal risk for Keyence is margin compression, not revenue collapse. The 41x forward multiple is paid almost entirely for the durability of the 50%+ operating margin and the high return on capital — so the structural threat is anything that erodes pricing power. Two transitions are the candidates. The first is general-purpose AI vision: if vision-language models running on commodity cameras and standard compute reach industrial-grade reliability, the specialized value that Keyence (and Cognex) charge for in machine vision gets commoditized, and the margin on that portion of the business compresses toward hardware economics. The second is the maturation of low-cost Asian competitors — Chinese sensor and vision vendors steadily climbing the capability curve and forcing Keyence to either cede the price-sensitive volume or defend it at lower margin.

Neither is imminent. Keyence's diversification across half-a-dozen product categories means no single transition breaks it the way a single technology shift could break a pure-play, and its direct-sales-plus-application-engineering model has historically been the hardest part for competitors to replicate — you can copy a sensor, you cannot easily copy a 20-year-trained direct-sales organization that captures solution value. Keyence's own roadmap (AI-enabled sensors, the VS-series integrated-AI vision systems) is a credible

response to the AI-vision threat — it is productizing AI rather than being displaced by it. But the constraint on the multiple is real: at 41x forward earnings the market assumes the 52% margin is permanent, and the terminal risk is precisely that it is not. The realistic terminal scenario is not a blow-up but a slow de-rating if margins drift from 52% toward, say, 40% over a decade.

Bull / Gap / Optionality (Photoncap framing)

1. The 52% operating margin is a moat with no industrial peer. FY2025 (ended March 2026) operating margin near 52%, with the March quarter at 54.3% ex-Cadenas amortization (company release, April 27, 2026) — roughly triple a typical automation peer and well above Cognex's low-20s. This is structural, cycle-tested pricing power, and it is the entire reason the premium multiple has held for two decades.

2. Five consecutive record years and accelerating overseas growth. FY2025 net sales of ¥1.16 trillion were a fifth straight record, and every overseas region grew 20%+ in the March 2026 quarter. The growth is broad and re-accelerating, not a single-cycle bounce, and overseas is now the larger and faster-growing portion of the base.

3. The most diversified customer and end-market profile in the batch. No dominant customer, no dominant end market, no dominant geography — spread across electronics, autos, pharma, logistics, food, semiconductors and general industrial worldwide. In a batch with a thin-margin turnaround and a 70%-auto cyclical, Keyence is the name nothing single can break.

4. Vision-guided robotics is a real, embedded growth leg. The VS series (integrated-AI vision-guided robotics) and CV-X 3D vision-guided robotics ride the structural shift to higher-mix, higher-speed factories. Keyence does not need to win any single robotics segment — it needs the automation adoption curve to keep rising and to keep capturing its share at its margin, which it has done through every cycle.

5. Fortress balance sheet and fabless model. Keyence holds a very large net-cash position and outsources most manufacturing — it carries minimal capital intensity, converts profit to cash efficiently, and has the financial resilience to invest through any downturn. It is the quality-compounder anchor of a robotics sensing portfolio.

Gap

1. The valuation prices permanence of a 52% margin. At 41.4x forward earnings and a ¥18.7-trillion market cap, the market assumes the extraordinary margin is

permanent. Any sustained drift in operating margin — from competition or AI-vision commoditization — would compress both earnings and the multiple simultaneously.

2. It is not a robotics pure-play and the theme exposure is undisclosed. Vision-guided robotics is embedded inside the machine-vision and sensor categories, not broken out as a segment. An investor buying the robotics theme through Keyence is really buying diversified factory automation — the robotics torque on the stock is muted versus a true pure-play.

3. The chart is moderately extended and the yen adds a layer. RSI 63.9 and +21.0% above the 50-day average is meaningfully extended (less than VPG, similar to Cognex). For a USD-based investor the Japan listing also adds JPY/USD translation risk and trading friction on top of the equity move.

4. Low-cost Asian competition is a slow but real margin threat. Chinese sensor and machine-vision vendors continue climbing the capability curve at far lower price points. Keyence can cede the price-sensitive volume or defend it at lower margin — either way, the pressure on the franchise's defining statistic builds over the back half of the decade.

Optionality

Event	Date / window	Direction
Q1 FY2026 (June quarter) earnings	~Aug 3, 2026	Binary on margin durability + overseas growth
Overseas regions sustain 20%+ growth into FY2026	FY2026 quarters	Bull — confirms broad acceleration
Operating margin holds ~52% vs. drifts lower	FY2026 reporting	Bear if margin slips
VS-series / AI-sensor adoption in robotics applications	2026–2027	Bull if it shows in mix
JPY/USD direction	Ongoing	Binary for USD-based holders

The trade

Keyence is the quality anchor of the sensing/vision batch — not the highest-torque robotics bet, but the one with the deepest moat and the lowest probability of a thesis-breaking surprise. The structural case: the most profitable sensor-and-vision company on earth, five straight record years, the most diversified customer base in the batch, with vision-guided robotics as a real embedded growth leg riding the factory-automation adoption curve. The trade is to initiate in a zone of current ±5%, roughly JPY 73,444–81,176, accumulating toward the lower end on any pullback; size at 0.75–1.0% of risk capital — a real but measured position, sized down from a true robotics pure-play because the theme exposure is diluted across the broader automation business and because the 41x multiple plus JPY translation risk argue for restraint. Stop at roughly JPY 64,000, below the structural base and the rising 50-day cloud; a break there would signal the margin-durability narrative is being questioned. The named catalyst is Q1 FY2026 (the June 2026 quarter) earnings in early August 2026, where the 52% operating margin and the 20%+ overseas growth get re-tested. For a USD-based investor who wants the machine-vision thesis without the Japan-listing friction, Cognex (CGNX, this batch) is the cleaner-access expression of the same vision-guided-robotics tailwind — though Keyence is the higher-quality business. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

MBLY — Mobileye Global Inc. · WATCH (Tier-2) · Conv 6/10 · Bucket B

Mobileye Global Inc. (MBLY)

The only profitable, self-funding name in AV perception — using its EyeQ ADAS cash engine to fund a robotaxi ramp and a $900M humanoid-robotics bet.

Investment Research · Photoncap-style deep dive · v1 of "Mobileye" · May 14, 2026

What Mobileye physically does

Mobileye sells the eyes and, increasingly, the driving brain for vehicles. The core product is the EyeQ system-on-chip — a custom automotive vision processor that takes raw camera (and increasingly radar and lidar) input and runs the computer-vision and neural-network workloads that detect lanes, vehicles, pedestrians, signs and free space. EyeQ ships inside the ADAS (advanced driver-assistance systems) modules of roughly 200 million vehicles cumulatively; it is the highest-volume dedicated automotive perception chip in the world, and the recurring per-unit royalty-plus-silicon revenue from it is the cash engine that funds everything else Mobileye does. EyeQ is the binding constraint in the sense that perception — knowing what is around the vehicle and where it is going — is the hard, safety-critical layer of any driver-assist or autonomous system, and Mobileye has two decades of road-validated data and a fleet-sourced mapping system (Road Experience Management, REM) that competitors cannot replicate quickly.

On top of the EyeQ base, Mobileye sells a tiered ladder of autonomy products. SuperVision is a hands-off, eyes-on "everywhere" point-to-point assisted-driving system; Chauffeur is the next rung, an eyes-off conditional-autonomy system; and Drive is the full robotaxi-grade stack Mobileye supplies to fleet operators. The strategic logic is that the same perception core, the same REM maps and the same compute family scale from a $1,300 SuperVision module up to a robotaxi brain — Mobileye monetizes the entire autonomy curve from one technology base rather than betting the company on robotaxi alone.

The 2026 strategic expansion is into physical AI beyond the car. In January 2026 Mobileye agreed to acquire Mentee Robotics — an AI-first humanoid-robotics company

founded by Mobileye's own CEO, Prof. Amnon Shashua — for roughly $900 million, a deal that closed in Q1 2026. The thesis Shashua articulated at CES 2026 is that the perception, world-modeling and sensor-fusion stack Mobileye built for vehicles is substantially transferable to humanoid and general-purpose robots, and that "physical AI" is the next addressable market for the same core competency. That is what puts MBLY squarely in the Robotics theme: not as a components supplier, but as a perception-and-autonomy platform deliberately extending from cars into robots.

Product roadmap

The EyeQ silicon roadmap runs through successive generations — EyeQ6 in volume now, with the high-end EyeQ6 High powering SuperVision and Chauffeur — and the company continues to win ADAS sockets across roughly three dozen OEMs. The dated, named product events that matter for the trade: SuperVision, priced around $1,300 per unit, is scheduled to integrate into Porsche and Audi models by late 2026; Chauffeur, priced between roughly $2,500 and $3,000 per unit, is set to launch with Audi in 2027, initially as an eyes-on system. On the design-win front, Mobileye disclosed at Q1 2026 a win with Mahindra covering SuperVision and Surround ADAS across at least six models from 2027 — a third Surround ADAS customer and a second SuperVision customer.

On robotaxi, the Drive program advances with named partners: Volkswagen's MOIA robotaxi program is moving toward series production on the Mobileye Drive stack, and Volkswagen Autonomous Mobility outlined at CES 2026 a roadmap targeting commercial robotaxi service across six cities by the end of 2027. Mobileye also continues robotaxi work with additional fleet partners. On the Mentee Robotics side, the disclosed roadmap is early: first on-site proof-of-concept humanoid deployments with customers are expected during 2026, intended to operate autonomously without teleoperation, with series production and commercialization targeted for 2028. What Mobileye does not do: it does not build cars, it does not operate robotaxi fleets itself (it supplies the stack to operators), and Mentee aside, it is not yet a robot manufacturer at scale — the humanoid business is a 2028-commercialization option, not a 2026 revenue line.

The financial print

Mobileye reported Q1 2026 results on April 23, 2026. Revenue was $558 million, up 27% year over year, driven by higher EyeQ unit volumes, increased ADAS fitment rates at core Western OEMs, and notably robust Chinese-OEM export demand. The headline

GAAP number was ugly and needs context: Mobileye took a roughly $3.8 billion goodwill impairment in the quarter, producing a GAAP net loss of about $3.8 billion. That impairment is a non-cash writedown of acquisition goodwill (largely the legacy Intel/Mobileye carrying value) — it does not touch the operating business or the cash. On an operating basis the quarter was strong: adjusted EPS rose roughly 51% year over year, and management raised the full-year 2026 revenue outlook to about $1.975 billion. The cash position is the differentiator in this batch — Mobileye is profitable on an adjusted basis, generates operating cash, and used roughly $591 million of cash to close the Mentee acquisition while still authorizing a $250 million share-repurchase program. Forward consensus puts the stock around a high-20s forward P/E (the ~29.2x in our data), which is a real earnings multiple — unlike every other name in this batch, MBLY is valued on profits, not on revenue hope. The next binary is Q2 2026 earnings, expected on or around July 23, 2026.

Customer mix today

Mobileye’s customer base is its structural strength: roughly three dozen OEM customers across the EyeQ ADAS franchise, with no single OEM dominant — there is no customer-concentration cliff of the kind that sinks single-program suppliers. The mix shift that matters in 2026 is geographic and product-tiered rather than customer-named. Geographically, Chinese-OEM export demand was explicitly called out as the swing factor in the Q1 2026 revenue beat — Chinese automakers exporting vehicles fitted with EyeQ, plus a normalization of Western dealer inventory after a soft 2024–25, drove the 27% growth. By product tier, legacy EyeQ ADAS remains the overwhelming majority of revenue and is the cash engine; SuperVision and Surround ADAS are the growth layer, still small in absolute revenue but with the Porsche, Audi and now Mahindra wins building a 2027-onward ramp; Drive (robotaxi) is pre-meaningful-revenue, monetizing through the VW/MOIA program toward 2027 series production; and Mentee/humanoid robotics is effectively $0 revenue today, a 2028-commercialization option. The honest framing: roughly 90%-plus of today’s revenue is the mature ADAS franchise, and the entire “robotics theme” case is the 2027–2028 layer on top.

What’s actually happening at the robotaxi and humanoid end-markets

The robotaxi mechanism is the VW/MOIA program. Mobileye supplies the Drive autonomous stack — perception, maps, driving policy — for Volkswagen's MOIA robotaxi service, which is advancing toward series production, and at CES 2026 VW Autonomous Mobility committed to a roadmap of commercial robotaxi service in six cities by end-2027. The significance is that this is a named OEM partner with a production timeline and a fleet operator behind it, not a science project — if it ships on schedule it converts Drive from a cost center into a per-vehicle revenue stream. Separately, SuperVision is the nearer-term proof point: management noted at Q1 2026 that SuperVision is demonstrating positive out-of-the-box US performance on unplanned routes and in challenging weather, which is the kind of generalization result that matters for the Chauffeur eyes-off launch with Audi in 2027.

On humanoid robotics, the Mentee mechanism is earlier and more speculative. Mentee builds an AI-first humanoid robot; Mobileye's pitch is that its vehicle-grade perception, world-modeling and sensor-fusion stack ports onto that platform, compressing Mentee's development timeline. The disclosed milestones: first on-site customer proof-of-concept deployments during 2026, running autonomously without teleoperation, with series production targeted for 2028. The skeptical read is that Mobileye paid $900 million for a pre-revenue startup founded by its own CEO — a related-party dynamic the 2026 proxy statement flagged — and the synergy claim, while plausible, is unproven. The bull read is that owning a humanoid platform gives Mobileye a captive design partner to prove the perception-stack-transfers thesis, and that 2028 commercialization is option value the market is barely paying for inside a 29x forward multiple.

The competitive threat / NVIDIA, Qualcomm and the in-sourcing OEMs

Mobileye's competitive set has three layers. The first is merchant silicon-and-stack rivals: NVIDIA's DRIVE platform is the most credible, pushing a high-compute, software-defined autonomy stack with strong developer mindshare, and Qualcomm's Snapdragon Ride is taking ADAS-and-beyond sockets with aggressive pricing. Both are larger, better-capitalized companies for whom automotive is one vertical among many. The second is OEM in-sourcing — the structural threat Mobileye has lived with for years. Tesla's vision-only FSD stack is the proof of concept that an OEM can build perception in-house and skip the merchant supplier entirely; several large OEMs have stood up internal autonomy teams. The third is the robotaxi-operator stack — Waymo most prominently — which competes for the eventual robotaxi market Mobileye's Drive program targets, though Waymo is an operator rather than a stack-supplier and the overlap is partial.

The competitive rebuttal that supports the bull case: Mobileye's EyeQ volume, its REM crowd-sourced maps, and two decades of validated road data are genuinely hard to replicate, and most OEMs do not have Tesla's appetite or scale to in-source — which is why Mobileye keeps winning sockets (Mahindra in Q1 2026) even as the in-sourcing narrative runs.

The terminal risk

The terminal risk for Mobileye is that the EyeQ ADAS franchise — the cash engine that funds the robotaxi ramp and the Mentee bet — erodes faster than the growth layers scale to replace it. The erosion vectors are real: OEM in-sourcing (the Tesla model), merchant competition from NVIDIA and Qualcomm taking next-generation sockets, and the long-run risk that "vision-only with enough compute" commoditizes the perception layer Mobileye sells. If EyeQ revenue and margin compress through the late 2020s while SuperVision/Chauffeur are still ramping and robotaxi/Mentee are still pre-scale, Mobileye becomes a melting ice cube funding two long-dated options. The hedge is that this is precisely why Mobileye is the safest name in this batch despite the risk: it is profitable and self-funding today, it has a $250 million buyback and a fortress balance sheet, and it has roughly three to four years of visible EyeQ cash flow to land the transition. The multiple you can pay is constrained by the in-sourcing risk; the survival of the company is not in question the way it is for the micro-caps in this batch.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. It is the only self-funding name in the batch. Mobileye is profitable on an adjusted basis, generated the cash to pay roughly $591 million for Mentee, and still authorized a $250 million buyback — all in a quarter when it also took a $3.8 billion non-cash goodwill writedown. Every other name in this perception batch is burning cash on a thin balance sheet; MBLY funds its robotics optionality out of operating cash flow.

2. The core business beat and the outlook was raised. Q1 2026 revenue of $558 million was up 27% year over year, adjusted EPS up roughly 51%, and management lifted full-year 2026 revenue guidance to about $1.975 billion. The ADAS franchise the bears keep writing off is currently accelerating, driven by Chinese-OEM export demand and Western inventory normalization.

3. The 2027 product ramp is contracted, not hoped-for. SuperVision integrates with Porsche and Audi by late 2026; Chauffeur launches with Audi in 2027; the Mahindra win adds SuperVision and Surround ADAS across six-plus models from 2027. These are named OEMs with priced products ($1,300 SuperVision, $2,500–$3,000 Chauffeur) and dated timelines — a visible step-up in revenue mix toward higher-ASP autonomy.

4. Robotaxi has a real production partner. The VW/MOIA Drive program is advancing toward series production, with VW Autonomous Mobility targeting commercial robotaxi service in six cities by end-2027. This is a funded OEM partner with a fleet operator and a timeline — the most credible robotaxi-supplier position outside of Waymo's in-house stack.

5. The Mentee humanoid bet is cheap option value inside the multiple. At ~29x forward earnings the market is paying for the ADAS-and-SuperVision business; the 2028-commercialization humanoid platform is barely priced. If the perception-stack-transfers thesis works even partially, Mentee is a free call option on physical AI bought with cash MBLY could spare.

Gap

1. ~90%-plus of revenue is still the mature ADAS franchise the bears want to short. The robotics theme case is the 2027–2028 layer; today MBLY is overwhelmingly a legacy EyeQ company. If in-sourcing and NVIDIA/Qualcomm competition bite before SuperVision and robotaxi scale, the cash engine compresses on schedule with the bear thesis.

2. The $3.8 billion goodwill impairment is a bad look even if non-cash. A write-down of that magnitude is the accountants confirming that the Intel-era carrying value was too high — it does not touch cash, but it is an admission that prior-regime expectations were not met, and it will be used against the stock.

3. The Mentee deal has a related-party stench. Mobileye paid roughly $900 million for a pre-revenue humanoid startup founded by its own CEO, Amnon Shashua — a conflict the 2026 proxy statement had to address. The synergy claim is plausible but unproven, and the price was set inside an obvious related-party dynamic.

4. RSI 77 and +30% above the 50-day MA — this is a momentum entry, not a value one. MBLY has run hard into the robotics narrative. At a high-20s forward multiple with the stock technically extended, a Q2 print that merely meets expectations could trigger a sharp mean-reversion; the entry has to respect that.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~Jul 23, 2026	Binary on FY26 ADAS ramp
SuperVision integration with Porsche / Audi	Late 2026	Bull if on schedule
First Mentee on-site humanoid POC deployments	During 2026	Bull
Chauffeur launch with Audi	2027	Bull
VW/MOIA robotaxi series production	Toward 2027	Bull
New ADAS socket losses to NVIDIA / Qualcomm	2026 ongoing	Bear

---

The trade

Mobileye is the Bucket-B anchor of the sensing/perception batch — the one name you can own without lying awake about the balance sheet, because it funds its robotics optionality out of operating cash flow. Initiate at $9.98–$11.04 (current price ± ~5%, the correct construction for a name running at RSI 77 well above its 50-day MA), size at roughly 1.5% of risk capital — a notch above the speculative micro-caps in this batch because MBLY is profitable and self-funding, but still measured given the in-sourcing overhang — and set the stop at $8.80, below the 50-day EMA cloud and the pre-breakout consolidation shelf. The named catalyst is Q2 2026 earnings around July 23, 2026, the binary on whether the Chinese-export-driven ADAS acceleration holds and the FY26 raise is conservative. The cleaner way to think about MBLY: it is not really a robotics pure-play yet — the Mentee humanoid layer is a 2028 option — so own it as the profitable AV-perception compounder with a free physical-AI call option attached, and accept that the robotics theme membership is mostly forward-looking. If you want unhedged robotics-sensing torque and can stomach the burn, OUST and AEVA give more beta to the theme; MBLY is the quality leg. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

OUST — Ouster, Inc. · WATCH (Tier-2) · Conv 6/10 · Bucket B

Ouster, Inc. (OUST)

The highest-unit-volume Western digital-lidar pure-play, repositioning from automotive supplier to a physical-AI sensing platform — revenue is inflecting but the balance sheet is thin.

Investment Research · Photoncap-style deep dive · v1 of "Ouster" · May 14, 2026

What Ouster physically does

Ouster builds digital lidar — laser sensors that measure distance by timing how long light pulses take to return, then assemble those returns into a 3D point cloud the host system uses to perceive its surroundings. The distinguishing word is "digital." Where legacy lidar uses discrete analog components — individual lasers, individual photodetectors, hand-aligned optics — Ouster's architecture puts the laser array and the detector array onto two custom CMOS chips. The light source is a VCSEL (vertical-cavity surface-emitting laser) array; the receiver is a SPAD (single-photon avalanche diode) array. Because the expensive, fiddly parts are silicon, the sensor rides the semiconductor cost-and-scaling curve rather than the optics-assembly curve. That is the entire investment premise: a sensor whose bill of materials and manufacturability improve like a chip, not like a camera module.

Functionally, the core product is the OS series — a family of spinning sensors (OS0 short-range wide-field, OS1 mid-range, OS2 long-range) that rotate to give a 360-degree field of view, plus the DF series of solid-state digital flash sensors aimed at automotive forward-facing applications. Each sensor is built around an Ouster-designed system-on-chip; the current generation is the L3 chip, and the just-launched Rev8 generation runs on next-generation "L4 Ouster Silicon." A lidar sensor is the binding constraint for any robot or vehicle that has to operate in unstructured 3D space in the dark, in glare, or where a camera's depth estimate is too noisy to trust — warehouses, mining sites, construction equipment, last-mile delivery robots, and the smart-infrastructure cameras that meter traffic and count people. Ouster's bet is that "physical AI" — the

current industry shorthand for robots and autonomous machines that act in the real world — needs a cheap, reliable, mass-manufacturable depth sensor, and that digital lidar is it.

Ouster reached its current scale partly through consolidation. It merged with Velodyne in 2023, absorbing the company that effectively invented automotive spinning lidar, and in early 2026 it folded in Stereolabs, a stereo-camera and spatial-perception company. The Stereolabs deal matters because it moves Ouster from “we sell you a lidar” toward “we sell you the perception stack” — lidar plus cameras plus the software that fuses them — which is a higher-value, stickier position if it executes.

Product roadmap

The legacy backbone is the OS sensor family — OS0, OS1, OS2 — which has shipped in successive revisions for years and remains the volume product, cumulatively well past 12,600 units in a single quarter as of Q1 2026. The DF series of solid-state digital flash sensors is the automotive forward-facing line, positioned against solid-state competitors but a smaller revenue contributor than the spinning OS family.

The headline 2026 event is the Rev8 OS family, released May 4, 2026 and billed by Ouster as “the world’s first native color lidar” — sensors that capture per-point color alongside range, eliminating a separate camera-to-lidar calibration step. Rev8 runs on the next-generation L4 Ouster Silicon and Ouster claims up to double the range and resolution of the prior generation, with an explicit emphasis on functional safety and automotive-grade reliability. Crucially for this theme, Ouster announced on May 5, 2026 that Rev8 is being brought to “the robotics and edge AI ecosystem” with native integration across the NVIDIA Jetson platform — the dominant compute module for autonomous mobile robots, drones, and cobots — and on May 12, 2026 that Rev8 has been qualified to run on NVIDIA’s DRIVE Hyperion platform for L4 autonomous-vehicle development. The NVIDIA tie-ins are the cleanest evidence that the robotics pivot is real and not just marketing.

What Ouster does not make: it does not make FMCW (frequency-modulated continuous-wave) lidar — the coherent, velocity-sensing approach that Aeva sells — nor does it make 4D imaging radar. It also does not make the host compute or the perception software’s final decision layer; with Stereolabs it now sells more of the perception stack, but it remains a sensor-and-middleware company, not an autonomy company.

The financial print

Ouster reported Q1 2026 results on May 5, 2026. Revenue was $48.6 million, up 49% year over year, with product revenue of $48.2 million on more than 12,600 lidar and camera units shipped. GAAP gross margin was 43%, down sharply from 60% in Q4 2025 — a swing management attributed to mix and the integration of lower-margin Stereolabs camera revenue, and a number to watch closely. The company posted a GAAP net loss of $17.5 million; adjusted EBITDA was a loss of $6.9 million, an improvement year over year but still not break-even. Management guided Q2 2026 revenue to $49.5–$52.5 million, including a full quarter of Stereolabs. Cash, cash equivalents and restricted cash stood at $80.5 million as of March 31, 2026 — a real number but not a large one for a company still burning, which makes the path to adjusted-EBITDA break-even the central financial question. There is no meaningful forward P/E; the company is loss-making, so the valuation rests on revenue growth and the gross-margin trajectory. The roughly $2.1 billion market cap implies the market is paying about 10x a run-rate revenue of ~$200 million for a sensor company that is not yet profitable — a momentum multiple, not a value one. The next binary is Q2 2026 earnings, expected on or around August 5, 2026.

Customer mix today

Ouster’s customer base is unusually fragmented for a company this size, which is both a strength (no single-customer cliff) and a weakness (no anchor program). Management does not break revenue out by named customer the way a single-OEM supplier would, and as of Q1 2026 reporting no single customer is disclosed as more than ~10% of revenue. By vertical, the structural shift is the story: in earlier years Ouster’s narrative was automotive-and-trucking heavy, leaning on the Velodyne legacy; by Q1 2026 management explicitly called out smart infrastructure and industrial automation as the key demand drivers of the quarter, citing new million-dollar contracts for the Ouster BlueCity smart-intersection product and several million-dollar industrial-automation deals. Robotics and edge AI is the fastest-growing slice — management flagged “strong demand from companies building foundational AI models and advanced robotics platforms” — but it is growing off a small base and is not yet separately quantified. Automotive is now the long-tail vertical rather than the headline. The Rev8 launch customer list Ouster published on May 5, 2026 is the best available proxy for direction of travel: it names Google, Volvo Autonomous Solutions, Liebherr, Epiroc, Field AI, Flyability, Skydio, Plus, Seegrid, Gecko Robotics, Burro, Third Wave Automation and roughly a dozen others — a mix heavily weighted toward robotics, mining, construction and warehouse automation, with autonomous passenger vehicles a minority. Treat that

list as design-in intent, not booked revenue.

What’s actually happening at the robotics / NVIDIA end-market

The most concrete development for the robotics thesis is Ouster’s two-step integration with NVIDIA. On May 5, 2026 Ouster announced that the Rev8 OS family is natively integrated across the NVIDIA Jetson platform — the Jetson Orin and Thor modules are the de facto compute standard for autonomous mobile robots, inspection drones and cobots, so being a plug-and-play sensor in that ecosystem lowers the integration cost for every robotics OEM building on Jetson. On May 12, 2026 Ouster added that Rev8 is qualified on NVIDIA DRIVE Hyperion, NVIDIA’s reference platform for L4 autonomous-vehicle development, with integration into NVIDIA DriveWorks. The mechanism that matters: NVIDIA does not pick a sensor lightly, and “qualified on the reference platform” means every robotaxi or AV developer prototyping on Hyperion can specify Ouster without a custom bring-up. The named robotics customers — Field AI (foundation models for robots), Skydio (autonomous drones), Gecko Robotics (industrial inspection robots), Seegrid and Third Wave Automation (warehouse AMRs and forklifts), Burro (agricultural robots) — are exactly the early-cycle physical-AI cohort the theme is built on. The honest caveat: these are design-ins and ecosystem placements, not multi-year production POs with disclosed unit volumes, and the robotics revenue line is still small enough that Ouster does not break it out.

The competitive threat / Hesai and RoboSense

The terminal competitive fact about lidar is that the volume and the cost curve now sit in China. Per industry tallies cited at CES 2026, Chinese players account for roughly 60% of automotive lidar revenue; Hesai shipped over 1.2 million units in 2025, reported full-year non-GAAP profitability, and told reporters at CES 2026 it would double production again on “accelerating demand” in automotive and robotics, with lidar now in roughly 25% of new EVs sold in China. RoboSense delivered on a similar order of magnitude. Against that, Ouster’s ~$200 million revenue run-rate and 12,600 units a quarter is a niche position. The bull rebuttal is that Hesai and RoboSense are automotive-volume players whose Western robotics and infrastructure penetration is limited by procurement caution and, increasingly, by US policy friction around Chinese sensors in critical infrastructure and defense-adjacent applications — which is precisely where Ouster’s BlueCity and industrial business sits. On the Western side, the competitive field has actually thinned: Luminar filed for Chapter 11 and is winding down or selling its lidar

business, removing a long-time bear-case comparison and a money-losing price competitor. Aeva (AEVA) competes on a different technology (FMCW) and a different segment (long-range automotive and industrial), so it is more an adjacent name than a head-to-head rival. The net read: Ouster is not the cost leader and never will be, so its survival case rests on owning the Western robotics/infrastructure niche where Chinese supply is structurally disadvantaged.

The terminal risk

Two structural transitions constrain the multiple. The first is the Chinese cost curve: if Hesai and RoboSense drive automotive-grade lidar ASPs low enough fast enough, and if Western procurement caution erodes, Ouster's addressable margin pool shrinks regardless of unit growth — you can win the robotics niche and still not earn an acceptable return on it. The second is sensor-architecture displacement: a meaningful camp in robotics argues that camera-only or camera-plus-radar perception, with enough neural-net compute, makes lidar optional for many indoor and structured-environment robots — the same "vision-only" argument Tesla made in automotive. If that view wins in warehouse AMRs and cobots, Ouster's fastest-growing vertical is also its most exposed. Ouster's hedge against both is real but unproven: the digital-silicon architecture is the credible answer to the cost curve (it scales like a chip), and the Stereolabs acquisition plus native-color Rev8 is the answer to vision-only (sell the fused lidar-plus-camera stack rather than fighting cameras). Whether those hedges are enough is exactly what the next two years of gross margin and robotics revenue disclosure will reveal.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. Revenue is genuinely inflecting, not just guided to inflect. Q1 2026 revenue of $48.6 million was up 49% year over year, and Q2 guidance of $49.5–$52.5 million keeps the run-rate above $200 million annualized. For a company that spent years stuck near break-even revenue scale, two consecutive quarters of ~50% growth driven by infrastructure and industrial demand — not a single fragile OEM program — is the most important fact in the file.

2. The NVIDIA double-integration is a real moat-builder. Native Jetson integration (May 5, 2026) plus DRIVE Hyperion qualification (May 12, 2026) means Ouster is

the default depth sensor in the two NVIDIA ecosystems that matter for robotics and AV development. Every Jetson-based robot OEM and every Hyperion-based AV developer can now specify Ouster without custom bring-up — that lowers Ouster's customer-acquisition cost across the entire physical-AI cohort.

3. The Western competitive field is thinning in Ouster's favor. Luminar's Chapter 11 removes a chronic Western price competitor and a bear-case comparison. With Hesai and RoboSense facing procurement and policy friction in US infrastructure, industrial and defense-adjacent applications, Ouster has a cleaner run at the exact niche — BlueCity smart infrastructure, mining, construction, warehouse robotics — where it is winning million-dollar contracts today.

4. Rev8 plus Stereolabs moves Ouster up the value stack. Native-color lidar eliminates a calibration step; the Stereolabs acquisition adds stereo cameras and fusion software. Together they let Ouster sell a perception bundle rather than a commodity sensor — a stickier, higher-ASP position if the integration delivers, and a direct answer to the "vision-only" bear case.

5. No customer-concentration cliff. With no disclosed customer above ~10% of revenue and demand spread across infrastructure, industrial, robotics and automotive, Ouster lacks the single-program risk that sinks most small sensor companies. The flip side is in the Gap section — but the diversification itself is a genuine durability feature.

Gap

1. The gross-margin reversal is alarming and unexplained-away. GAAP gross margin fell from 60% in Q4 2025 to 43% in Q1 2026. Management attributed it to mix and Stereolabs integration, but a 17-point swing in one quarter undermines the entire "scales like a chip" thesis if it is not promptly reversed. Until margin re-expands, the digital-lidar cost-curve story is a claim, not a result.

2. The balance sheet is thin for a company still burning. Cash of $80.5 million against an adjusted-EBITDA loss of $6.9 million per quarter and a GAAP net loss of $17.5 million is a runway measured in quarters, not years, on GAAP terms. Any stumble in the revenue ramp or the margin recovery raises the specter of a dilutive raise at a momentum-inflated share count.

3. Robotics revenue is real in narrative but invisible in the financials. Management talks up robotics and edge AI as the fastest-growing vertical, but does not quantify it. The investment case for this name within the Robotics theme is therefore largely option value — design-ins and ecosystem placements that have not yet shown up as a disclosed, growing revenue line.

4. Ouster will never be the cost leader. Hesai shipped 1.2 million units in 2025 and is

profitable; RoboSense is at similar scale. Ouster ships ~12,600 a quarter. The structural cost and scale advantage sits in China, and Ouster's entire survival case depends on a policy-and-procurement moat around its Western niche holding — a moat that is real today but not guaranteed to last.

Optionality

Event	Date / window	Direction
Q2 2026 earnings — gross-margin recovery test	~Aug 5, 2026	Binary on the cost-curve thesis
First disclosed/quantified robotics revenue line	H2 2026	Bull if broken out and growing
Rev8 design-ins converting to production POs	H2 2026 – 2027	Bull
Further NVIDIA / hyperscaler robotics partnerships	2026 ongoing	Bull
Capital raise / dilution on thin cash	If margin or ramp slips	Bear

---

The trade

Ouster is a Bucket-B, early-cycle option on Western digital lidar becoming the default depth sensor for physical AI — a real revenue inflection wrapped around a thin balance sheet and a single ugly margin print. Initiate at $31.60–$34.90 (current price ± ~5%, the appropriate construction for a name that just gapped up on the NVIDIA news and is mid-melt-up), size at roughly 1.0% of risk capital — speculative tier, because this is a loss-maker whose robotics revenue is still narrative — and set the stop at $26.50, below the post-Rev8 breakout shelf and the rising 50-day EMA cloud. The named catalyst is Q2 2026 earnings around August 5, 2026, which is the binary on whether the 43% gross margin was a one-quarter Stereolabs-integration artifact or a structural problem; a clean margin recovery plus continued ~50% revenue growth would re-rate the name, a second soft margin print would break the thesis. If you want the same physical-AI-sensing exposure with a fortress balance sheet and actual profits, Mobileye (MBLY) is the cleaner expression — slower-growth but self-funding — and serious lidar-volume

bulls should note Hesai is the profitable scale player, albeit with China-risk. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

AEVA — Aeva Technologies, Inc. · WATCH (Tier-2) · Conv 5/10 · Bucket B

Aeva Technologies, Inc. (AEVA)

The leading 4D FMCW lidar-on-chip — genuinely differentiated technology with marquee design wins, but revenue is tiny and the robotics case is still pure option value.

Investment Research · Photoncap-style deep dive · v1 of "Aeva" · May 14, 2026

What Aeva physically does

Aeva builds a fundamentally different kind of lidar. Most lidar — including Ouster's, Hesai's, RoboSense's — is time-of-flight: it fires a pulse of light and times how long the reflection takes to return, which gives distance. Aeva's sensors are FMCW — frequency-modulated continuous-wave. Instead of a pulse, an FMCW lidar emits a continuous laser beam whose frequency is swept up and down, then mixes the returning light with the outgoing beam. Because the beam is coherent, the interference pattern encodes not just range but also velocity directly — the Doppler shift of the returning light tells the sensor how fast each individual point is moving toward or away from it, instantly, per point, on the first frame. That is what "4D" means: x, y, z, plus velocity. A time-of-flight lidar has to see an object across multiple frames and infer its motion; an Aeva sensor measures the motion of every pixel in the point cloud natively. For a vehicle or a robot trying to distinguish a pedestrian stepping off a curb from a parked one, or a robot arm tracking a moving part on a conveyor, instant per-point velocity is a real perceptual advantage.

The second differentiator is integration. Aeva calls its product "4D LiDAR-on-chip" — the photonics that would normally be discrete optical components are integrated onto a silicon photonics module, the same scale-and-cost logic Ouster applies to detectors but applied to the coherent FMCW front end. FMCW is also inherently more interference-resistant: because each sensor only detects light coherent with its own swept beam, it does not get blinded by sunlight or by other lidars, a growing problem as lidar density rises. The product families: the Atlas automotive 4D sensor, the new Atlas Ultra (a slimmer, behind-the-windshield-capable variant shown at CES 2026), and Eve, the tech-

nology variant aimed at precision industrial and metrology applications. The binding-constraint case for Aeva is narrower than for a volume lidar player: it is not “cheapest depth sensor for every robot,” it is “the highest-fidelity, velocity-aware, interference-immune sensor for applications where perception quality is safety-critical or precision-critical” — autonomous trucks, premium passenger AVs, and industrial inspection.

Product roadmap

The automotive line is the Atlas family. Atlas is Aeva’s production-intent 4D automotive sensor; in Q1 2026 Aeva delivered production-intent Atlas sensors to Daimler Truck and disclosed it had delivered initial Atlas 4D C-sample units for Daimler Truck North America and Torc Robotics’ planned Level 4 autonomous Freightliner Cascadia program. Aeva is the exclusive long-range lidar supplier for Daimler Truck’s autonomous-truck production program, with start of production planned for 2026 and the autonomous-truck market launch targeted around 2027. Atlas Ultra — a slimmer sensor designed to mount behind a passenger vehicle’s windshield, co-developed with a purpose-built windshield from Wideye by AGC — debuted at CES 2026 and is the product behind Aeva’s passenger-OEM ambitions. On the passenger side, Aeva disclosed that a global top-10 passenger OEM awarded it a development program for that OEM’s next-generation global vehicle platform, with a stated letter of intent toward a large-scale multi-year production program award “this year” — i.e., a 2026 conversion event, not yet a signed production contract.

The industrial line is Eve. Aeva’s Eve technology powers Nikon’s new APDIS MV5X Laser Radar system for automated industrial inspection — a shipping, named-customer product — and Aeva has named SICK AG and LMI Technologies as additional industrial-automation partners. At CES 2026 Aeva also introduced new sensor technology explicitly positioned for “physical AI,” its framing for the robotics opportunity. What Aeva does not make: it does not make time-of-flight lidar, it does not make radar or cameras, and it is not a perception-software or autonomy company — it is a sensor supplier whose differentiation is the FMCW physics. It also does not yet make anything at production volume; every program above is C-sample, development-award or NRE stage.

The financial print

Aeva reported Q1 2026 results on May 7, 2026. Revenue was $6.3 million, up roughly 90% year over year from $3.4 million — but the growth was driven mainly by a sharp

increase in non-recurring engineering (NRE) services revenue, not product shipments, which is the single most important caveat in the file: this is engineering-services revenue from development programs, not a product run-rate. The net loss was $35.0 million, roughly flat year over year, on continued heavy R&D spending and elevated stock-based compensation. Operating activities used $25.8 million of cash in the quarter. The balance sheet is the relative bright spot: cash, cash equivalents and marketable securities totaled $99.5 million, and Aeva disclosed total available liquidity of about $224.5 million including a fully undrawn $125 million credit facility. At roughly $25.8 million of quarterly operating burn, the $99.5 million of on-balance-sheet cash is a runway of roughly four quarters before drawing the facility — adequate, not comfortable, and a number that makes the timing of the passenger-OEM production conversion genuinely matter. There is no forward P/E; Aeva is deeply loss-making and the ~$1.34 billion market cap is roughly 50x trailing revenue — a pure technology-and-design-win multiple. The next binary is Q2 2026 earnings, expected on or around August 6, 2026.

Customer mix today

Aeva's revenue is too small and too NRE-heavy to have a conventional production-volume customer mix — the more honest description is a portfolio of development programs at different maturities. The anchor is Daimler Truck: Aeva is the exclusive long-range lidar supplier for Daimler Truck's autonomous-truck production program, delivered production-intent Atlas sensors and C-samples for the Daimler Truck North America / Torc Robotics autonomous Freightliner Cascadia program in Q1 2026, with start of production planned for 2026 — this is the program most likely to become real production revenue first. The passenger-vehicle leg is a global top-10 OEM development award for that OEM's next-generation platform, with a letter of intent toward a multi-year production award targeted for 2026 — a high-value option, not yet revenue. The industrial leg is the most commercially mature in terms of actually-shipping product: Aeva's Eve technology is inside Nikon's APDIS MV5X laser-radar inspection system, a named customer with a shipping product, with SICK AG and LMI Technologies as additional industrial partners. The structural shift the company wants investors to see: from a 2024 narrative that was almost entirely "automotive development programs" to a 2026 narrative of "diversified across trucking, passenger and industrial" — but the financial reality is that none of these is yet a disclosed production-volume revenue line, and the 90% revenue growth was NRE, not units.

What’s actually happening at Daimler Truck and the industrial end-market

The Daimler Truck mechanism is the closest thing Aeva has to a visible path to real revenue. As the exclusive long-range lidar supplier for Daimler’s autonomous-truck program, Aeva moved in Q1 2026 from earlier-stage samples to production-intent Atlas sensors and C-sample units for the Daimler Truck North America / Torc Robotics L4 Freightliner Cascadia program. The significance: “production-intent” and “C-sample” are the late-stage qualification rungs before a sensor goes into series production, and Daimler has stated a 2026 start-of-production target with an autonomous-truck market launch around 2027. If that timeline holds, Daimler converts from an NRE customer into a per-unit product customer — the inflection the whole bull case needs. The skeptical note: autonomous-trucking start-of-production timelines have a long history of slipping across the entire industry, and Aeva does not control Daimler’s schedule.

The industrial mechanism is smaller but already real. Aeva’s Eve technology is shipping inside Nikon’s APDIS MV5X laser-radar system for automated industrial inspection — this is a named OEM, a shipping product, and a vertical (precision metrology and inspection) where FMCW’s per-point velocity and interference immunity are genuine advantages and where the customer is not price-shopping against $200 Chinese time-of-flight units. SICK AG and LMI Technologies extend the same logic into factory automation. This industrial leg is where the robotics theme membership is most tangible today: it is sensing for industrial machines and inspection robots, shipping now, even if it is a small share of a small revenue base. The CES 2026 “physical AI” sensor introduction is the explicit signal that Aeva intends to push Eve-class technology into general robotics — but that is roadmap, not revenue.

The competitive threat / Chinese time-of-flight scale and Ouster

Aeva competes on two fronts. The first is the broad lidar field, where the structural fact is Chinese scale: Hesai shipped over 1.2 million units in 2025 and is profitable, RoboSense is at similar scale, and Chinese players hold roughly 60% of automotive lidar revenue per CES 2026 industry tallies. Aeva’s defense is that it is not trying to win the volume time-of-flight market — it is selling a higher-fidelity, velocity-native, interference-immune sensor at a premium for safety-critical and precision-critical applications. That is a real differentiation, but it is also a structurally smaller addressable pool, and it means Aeva can be technologically right and still fail to reach profitable volume. The second front is other FMCW and Western players: Aeva is the most prominent pure-play FMCW lidar company, which is an advantage (it owns the category narrative) but also

a warning (FMCW has been “about to win” for years without displacing time-of-flight at scale). Among Western peers, Luminar’s Chapter 11 filing removed a long-range-lidar competitor and a chronic price-war participant; Ouster (OUST) competes in robotics and infrastructure but on time-of-flight digital lidar, so it is more an adjacent name than a head-to-head rival. The cleanest competitive summary: Aeva’s technology is genuinely differentiated, but differentiation has not yet translated into the production volume that would prove the FMCW cost curve can close.

The terminal risk

The terminal risk for Aeva is that the FMCW cost-and-complexity premium never closes. FMCW lidar is harder to build than time-of-flight — coherent detection, swept-frequency lasers, silicon photonics integration — and it has always carried a cost and yield disadvantage. The bull thesis depends on Aeva’s lidar-on-chip integration driving that cost down the silicon curve toward time-of-flight parity while keeping the velocity-and-interference advantages. If that convergence does not happen — if FMCW stays meaningfully more expensive while Chinese time-of-flight lidar keeps getting cheaper and “good enough” for most applications — then Aeva is permanently a niche performance supplier serving autonomous trucks, premium AVs and industrial inspection, a real but small market that may never support a profitable company at Aeva’s cost structure. The named alternative beneficiaries are the time-of-flight scale players (Hesai, RoboSense) and, in robotics specifically, camera-and-radar fusion approaches that skip high-end lidar entirely. Aeva has a credible roadmap argument — the on-chip integration — but it is unproven at volume, and that uncertainty is exactly what caps the multiple on a company this far from profitability.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. The technology is genuinely differentiated, not marketing. FMCW measures per-point velocity natively and is immune to sun and cross-lidar interference — real perceptual advantages for safety-critical and precision applications. Aeva is the leading pure-play in the category, and as lidar density rises, interference immunity becomes a more valuable property, not less.

2. Daimler Truck is a marquee, late-stage program. Aeva is the exclusive long-range lidar supplier for Daimler’s autonomous-truck program and delivered production-

intent Atlas sensors plus C-samples for the Daimler Truck North America / Torc L4 Cascadia program in Q1 2026, with 2026 start-of-production targeted. This is the most credible path to converting NRE into product revenue.

3. The industrial leg is shipping product now. Aeva's Eve technology is inside Nikon's APDIS MV5X inspection system — a named OEM, a shipping product — with SICK AG and LMI Technologies extending into factory automation. This is real robotics-adjacent revenue today, in a vertical where Aeva is not price-shopped against commodity lidar.

4. A top-10 passenger OEM development award with a 2026 LOI conversion. A global top-10 passenger automaker awarded Aeva a development program for its next-generation global platform, with a stated letter of intent toward a large-scale multi-year production award targeted for 2026. If that converts, it is the single most material catalyst in the file.

5. Liquidity is adequate and the field is thinning. $99.5 million of cash plus a $125 million undrawn facility gives Aeva room to reach its 2026 conversion milestones, and Luminar's Chapter 11 removes a long-range-lidar competitor and a price-war participant.

Gap

1. The revenue is tiny and the "90% growth" was NRE, not product. Q1 2026 revenue of $6.3 million grew mainly on non-recurring engineering services — development-program fees, not a product run-rate. At ~50x trailing revenue, the valuation is paying for design wins, not a business.

2. The loss is deep and the runway is finite. A $35.0 million quarterly net loss and $25.8 million of operating cash burn against $99.5 million of on-balance-sheet cash is roughly four quarters before drawing the credit facility. The passenger-OEM conversion has to land on schedule, or a dilutive raise looms.

3. Every flagship program is pre-production. Daimler is C-sample, the passenger OEM is a development award, the industrial wins are small. Nothing in the file is yet a disclosed production-volume revenue line, and autonomous-trucking start-of-production timelines slip across the whole industry.

4. FMCW has been "about to win" for years. The category's cost-and-yield disadvantage versus time-of-flight is structural and persistent. If Aeva's on-chip integration does not close that gap while Chinese time-of-flight keeps getting cheaper, Aeva is technologically right and commercially stuck.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~Aug 6, 2026	Binary on NRE-to-product transition
Top-10 passenger OEM LOI converts to production award	During 2026	Bull — most material catalyst
Daimler Truck start of production	2026	Bull if on schedule
Atlas Ultra passenger-OEM design wins	2026–2027	Bull
New “physical AI” / robotics sensor design-ins	2026 ongoing	Bull
Capital raise / credit-facility draw	If conversion slips	Bear

---

The trade

Aeva is a Bucket-B speculative option on FMCW lidar proving its cost curve can close — genuinely differentiated technology and marquee design wins wrapped around a tiny, NRE-heavy revenue base and a four-quarter cash runway. Initiate at $20.16–$22.28 (current price ± ~5%; note one secondary source showed AEVA near $19 on May 12, a roughly 9–10% gap to the reference price worth verifying before sizing up), keep the size small at roughly 0.75% of risk capital — the deepest-loss-maker in this batch alongside the tiny revenue base argues for the lightest weight — and set the stop at $16.50, below the pre-rally consolidation shelf. The named catalyst is Q2 2026 earnings around August 6, 2026, but the real binary is whether the top-10 passenger-OEM letter of intent converts to a signed production award during 2026; that conversion, more than any quarterly print, is the event that re-rates the name. If you want differentiated lidar exposure with a far better balance sheet and an actual revenue inflection, Ouster (OUST) is the cleaner expression of the same Western-lidar-into-robotics thesis; Aeva is the higher-beta, higher-technology-risk version. Be honest with yourself that this is option value on a loss-making balance sheet, not an investment in a business. **Conviction: 5 / 10.**

Sources referenced inline throughout. Reference v1 of this template format:

_Watchlist/hanmi-photoncap-style.md.

INDI — indie Semiconductor, Inc. · WATCH (Tier-2) · Conv 5/10 · Bucket B

indie Semiconductor, Inc. (INDI)

An automotive multimodal-sensing SoC company with real radar, vision and lidar revenue — and a genuine embodied-AI design pipeline sitting on a loss-making, debt-heavy balance sheet.

Investment Research · Photoncap-style deep dive · v1 of "indie Semiconductor" · May 14, 2026

What indie Semiconductor physically does

indie Semiconductor designs the mixed-signal and system-on-chip silicon that sits behind a vehicle's sensors. Where Mobileye sells the vision processor and Arbe sells the radar chipset, indie's franchise is broader and more horizontal: it makes the SoCs that drive radar, vision/camera, lidar and ultrasound sensing, plus connectivity and power chips for the broader vehicle. The company's positioning word is "multimodal" — the bet that a modern vehicle (and, increasingly, a robot) does not rely on a single sensor type but fuses several, and that a supplier with credible silicon across radar, vision, lidar and ultrasound can win more content per platform than a single-modality specialist. indie's products are the analog-and-digital building blocks: a radar SoC that handles the RF front end and signal processing, a vision processor that ingests camera data, an FMCW lidar SoC that runs the laser and detector control, and ultrasound chips for low-speed parking and proximity sensing.

The technically interesting product in indie's lineup is the iND83301 FMCW lidar SoC. indie's claim is that this chip enables roughly an 80% reduction in power consumption and a 40% reduction in solution size for an FMCW lidar system — meaningful, because power and size are exactly the constraints that keep high-end lidar out of battery-powered robots and compact machines. The other notable part is the iND880 vision processor, which indie has confirmed in production at leading EV manufacturers. The binding-constraint argument for indie is less dramatic than for a pure-play sensor company: indie is not the sensor, it is the silicon plumbing that makes a multimodal sensing system

buildable, cheap and power-efficient. What puts indie squarely in the Robotics theme is that management has been explicit that this same multimodal-sensing silicon — and specifically the iND83301 lidar SoC and the radar and vision SoCs — is being designed into autonomous mobile robots, humanoids and drones, not just cars. indie frames this as “embodied AI,” and unlike most names making that claim, indie has at least one disclosed robotics design-in to point to.

Product roadmap

indie’s roadmap is a portfolio of automotive sensing and connectivity SoCs rather than a single hero product with launch dates, but the named parts that matter are specific. The iND83301 FMCW lidar SoC is the embodied-AI flagship — indie disclosed it is now “adopted beyond automotive in autonomous robotics and embodied AI applications,” and management cited a specific win: indie’s lidar SoC for an Advanced Mobile Robot (AMR) at a major global logistics company. The iND880 vision processor is confirmed in production at leading EV manufacturers and is part of a pipeline indie describes as “tens of millions of annual revenue dollars” in opportunities, with management stating iND880-and-related vision revenue has the potential to exceed radar revenue in 2026. On radar, indie disclosed a $25 million production order from its Tier-1 radar partner, driven by demand from two automotive OEMs across European and Asian markets. Beyond sensing, indie has pushed into adjacent photonics: at Q1 2026 it introduced what it called “the world’s first commercially available ultraviolet distributed-feedback (DFB) laser at 399 nm,” a wavelength matched to the ytterbium cooling transition used in neutral-atom quantum-computing architectures — a niche but real product diversification.

What indie does not make: it does not make finished sensors or modules, it does not make the host autonomy compute, and it is not a robot or vehicle company — it is a fabless SoC supplier that reaches OEMs largely through Tier-1 integrators. It also does not break out a clean robotics revenue line; “embodied AI” is a disclosed and growing design pipeline, not yet a quantified segment.

The financial print

indie reported Q1 2026 results on May 7, 2026. Revenue was $55.5 million, up about 3% year over year and above the guidance midpoint — modest growth, reflecting an automotive-semiconductor end-market that has been soft. The non-GAAP operating loss was $11.1 million, an improvement from a $15.1 million non-GAAP operating loss in Q1

2025, so the loss is narrowing but the company is still not profitable. Management guided Q2 2026 revenue to $59–65 million ($62 million midpoint), of which roughly $25 million is the Wuxi business being divested and roughly $37 million is the core business. The balance sheet is the central concern: cash, cash equivalents and restricted cash were about $184.7 million, but total debt rose to roughly $415–428 million. During the quarter indie issued $170.5 million of 4.00% convertible notes due 2031 and used part of the proceeds to repurchase $104.0 million of its 4.50% notes due 2027 — a refinancing that pushed maturities out roughly four years and lowered the coupon, which is genuinely constructive, but it leaves indie with a debt load several times its cash and well above its annual revenue. A mitigant is the pending Wuxi divestiture, which indie expects to close later in 2026 for roughly $135 million of net cash proceeds. The forward P/E of roughly 31x reflects consensus expectations that indie reaches profitability — so unlike the pre-revenue names in this batch, the valuation embeds an actual earnings path, but it is a forward path, not a present one. The next binary is Q2 2026 earnings, expected on or around August 6, 2026.

Customer mix today

indie sells largely through Tier-1 integrators into a broad automotive-OEM base, so it does not have the single-customer concentration of a direct OEM supplier — but it also does not disclose a clean named-customer percentage breakout. The mix is better described by product line and end-market. On radar, the disclosed anchor is a $25 million production order from indie's Tier-1 radar partner, driven by demand from two automotive OEMs spanning European and Asian markets — a real, sized order from a named-category partner. On vision, the iND880 vision processor is confirmed in production at "leading EV manufacturers" (unnamed), and management framed vision as a "tens of millions of dollars" pipeline that could exceed radar revenue in 2026 — a notable structural shift, because it implies vision overtaking radar as indie's largest sensing line within the year. On lidar and embodied AI, the disclosed win is indie's lidar SoC for an Advanced Mobile Robot at a major global logistics company — the single most concrete robotics customer reference in the file, though unnamed and unquantified. The structural shift the company wants investors to see: from a 2024 mix that was almost entirely automotive across radar, vision and connectivity, toward a 2026 mix where vision is overtaking radar and a new embodied-AI/robotics line is appearing — a logistics AMR customer, plus disclosed engagement "across embodied AI applications, including autonomous mobile robots, humanoids and drones." The honest caveat: the robotics revenue is real enough to name a customer but small enough that indie does not size it, and the overwhelming majority of the $55.5 million quarterly revenue is still automotive.

What’s actually happening at the embodied-AI / robotics end-market

The mechanism that puts indie in this theme — and distinguishes it from names where “robotics” is pure marketing — is that indie’s automotive sensing silicon is power-efficient and compact enough to be designed directly into robots. The clearest example is the iND83301 FMCW lidar SoC: indie’s disclosed claim of roughly 80% lower power and 40% smaller solution size is precisely what an FMCW lidar needs to fit on a battery-powered AMR or a humanoid, and indie disclosed that the chip is now adopted beyond automotive specifically in autonomous robotics and embodied-AI applications, with a named (if unspecified) win — a lidar SoC for an Advanced Mobile Robot at a major global logistics company. Management described engagement “accelerating across embodied AI applications, including autonomous mobile robots, humanoids and drones.” The mechanism is credible because it is the same silicon, re-targeted: indie does not have to invent a robotics product, it has to win designs for chips it already ships into cars. The skeptical read is the one that applies to every name in this batch — the embodied-AI revenue is a design pipeline, not a quantified segment, and one logistics-AMR win does not move a $55 million quarter. The bull read is that indie’s robotics positioning is grounded in shipping silicon and at least one real customer, which is more than most “robotics-adjacent” semiconductor companies can say, and that the multimodal portfolio (radar + vision + lidar + ultrasound) is exactly what a robot’s sensor suite needs.

The competitive threat / NXP, Infineon, TI and Analog Devices

indie’s competitive problem is scale. It competes in automotive sensing and mixed-signal silicon against NXP Semiconductors, Infineon, Texas Instruments and Analog Devices — companies with tens of billions of dollars in revenue, decades of automotive-qualified track record, and the engineering budgets to iterate every modality indie touches. In radar specifically, NXP, Infineon and TI are the entrenched incumbents; in vision and lidar SoCs, indie faces both those incumbents and specialists. indie’s strategy against this is to be the nimble multimodal challenger — to win content by being faster, more integration-friendly, and willing to support smaller and newer programs (including robotics) that the giants deprioritize. That is a viable niche strategy, but it is a niche strategy: indie’s roughly $220 million annual revenue run-rate is a rounding error against its competitors, and in a soft automotive-semiconductor cycle the in-

cumbents can compress pricing in ways a sub-scale, loss-making, debt-laden challenger cannot easily absorb. The competitive question that decides indie's fate is whether the embodied-AI/robotics adjacency grows fast enough to give indie a higher-growth, less-contested revenue pool before the automotive core gets margin-squeezed.

The terminal risk

The terminal risk for indie is being permanently sub-scale: a loss-making, debt-heavy challenger in a market where the economics favor giants, with a robotics adjacency that stays too small to change the trajectory. The automotive-sensing-silicon market is structurally one where NXP, Infineon, TI and ADI have cost, scale and qualification advantages; if the embodied-AI revenue indie is counting on remains a design pipeline rather than a material segment through the late 2020s, indie is left as a sub-scale automotive supplier carrying a debt load several times its cash, dependent on a soft end-market recovering and on the Wuxi divestiture proceeds to manage liquidity. The named alternative beneficiaries are precisely the incumbents. indie's hedge is real but unproven: the multimodal portfolio plus the explicit, shipping-silicon push into robotics is a credible attempt to find a higher-growth, less-contested pool — and the debt refinancing bought time. But the multiple you can pay is constrained by the fact that the bull case requires both an automotive-cycle recovery and a robotics ramp, and the balance sheet does not give indie unlimited quarters to wait for both.

Bull / Gap / Optionality (Photoncap framing)

1. The robotics positioning is grounded in shipping silicon, not slides. indie's iND83301 lidar SoC — with disclosed ~80% lower power and ~40% smaller footprint — is adopted in autonomous robotics and embodied AI, with a named win (a lidar SoC for an AMR at a major global logistics company). Among "robotics-adjacent" semiconductor names, indie is one of the few with an actual robotics customer to point to.

2. The loss is narrowing and the revenue mix is upgrading. Non-GAAP operating loss improved to $11.1 million from $15.1 million a year earlier, and management expects vision (iND880) revenue to potentially exceed radar revenue in 2026 — a shift toward a higher-value, EV-driven product line, plus a sized $25 million Tier-1 radar order across European and Asian OEMs.

3. The debt refinancing materially de-risked the balance sheet. Issuing $170.5 million of 4.00% notes due 2031 and retiring $104.0 million of 4.50% notes due 2027

pushed maturities out roughly four years and cut the coupon — indie removed the near-term refinancing cliff that would otherwise hang over a loss-making company.

4. The Wuxi divestiture brings ~$135 million of net cash. Closing expected later in 2026, the divestiture would add roughly $135 million of net cash proceeds — a real liquidity cushion against the debt load and a signal management is willing to streamline the portfolio toward the core sensing franchise.

5. The forward multiple embeds an actual earnings path. Unlike the pre-revenue names in this batch, indie's ~31x forward P/E reflects consensus that it reaches profitability — the loss trajectory, the mix upgrade and the cost discipline together give that path more credibility than a pure hope multiple.

Gap

1. The debt load is the dominant risk. Roughly $415–428 million of total debt against ~$185 million of cash, on a company still posting operating losses. The refinancing bought time and the Wuxi proceeds will help, but indie is a leveraged loss-maker, and leverage plus losses is the combination that ends companies in a downturn.

2. Core revenue growth is anemic. Q1 2026 revenue grew only ~3% year over year. The automotive-semiconductor end-market is soft, and indie is a sub-scale player in it — the bull case needs both a cycle recovery and a robotics ramp, and neither is in hand.

3. The robotics revenue is unsized. "Embodied AI" is a disclosed and accelerating design pipeline with one named AMR customer — but indie does not break out a robotics revenue number, which means the theme membership is, as with the rest of this batch, largely option value rather than a current earnings driver.

4. The competitive set is overwhelming. NXP, Infineon, TI and ADI dwarf indie in scale, automotive-qualification track record and engineering budget. In a soft cycle the incumbents can compress pricing in ways a sub-scale, debt-laden challenger cannot absorb.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~Aug 6, 2026	Binary on core-revenue and lo
Wuxi divestiture close	Later in 2026	Bull — ~$135M net cash
Vision (iND880) revenue exceeding radar	During 2026	Bull if confirmed
New embodied-AI / humanoid / drone design wins	2026 ongoing	Bull
Automotive-semiconductor cycle recovery	2026–2027	Bull

Event	Date / window	Direction
Liquidity stress if cycle stays soft	2026–2027	Bear

The trade

indie is a Bucket-B speculative name — a real multimodal automotive-sensing SoC business with a genuine, shipping-silicon push into robotics, weighed down by a debt load that is the single biggest thing standing between the bull case and the bear case. Initiate at $4.25–$4.69 (current price ± ~5%), size at roughly 0.75% of risk capital — speculative tier, because the leverage-plus-losses combination demands respect even though indie has more of a real business than the pre-revenue names in this batch — and set the stop at $3.55, below the pre-rally base. The named catalyst is Q2 2026 earnings around August 6, 2026, but the cleaner binary to watch is the Wuxi divestiture close later in 2026, which converts a portfolio question into roughly $135 million of balance-sheet relief. The honest framing: indie's robotics revenue is nascent and unsized, so the embodied-AI theme membership is option value layered on an automotive-sensing turnaround that itself depends on a soft cycle recovering — two bets stacked on a leveraged balance sheet. If you want multimodal-sensing exposure without the debt overhang, Mobileye (MBLY) is the self-funding quality expression and Ouster (OUST) is the cleaner balance-sheet pure-play on sensing-into-robotics; indie is the higher-leverage, higher-torque version for an investor who believes both the automotive cycle and the embodied-AI ramp come through. **Conviction: 5 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

ST — Sensata Technologies Holding plc · WATCH (Tier-2) · Conv 5/10 · Bucket B

Sensata Technologies Holding plc (ST)

A cheap, mid-turnaround auto-and-industrial sensor company with a real force/position-sensing franchise — robotics is a thin option, not the story.

Investment Research · Photoncap-style deep dive · v1 of "Sensata Technologies" · May 14, 2026

What Sensata Technologies physically does

Sensata makes the sensors and electrical-protection components that let vehicles and industrial systems measure and protect themselves. The product catalogue is built on three physical-measurement domains: pressure (the company's heritage strength — pressure sensors for engines, brakes, HVAC, refrigeration and industrial systems), temperature, and force/position (sensors that measure mechanical force, torque and the position of moving parts), plus a large electrical-protection franchise (circuit breakers, contactors, fuses and high-voltage protection for electric vehicles). Sensata reorganized into three operating segments — Automotive; Industrials; and Aerospace, Defense and Commercial Equipment — and the common thread across all of it is taking a physical variable and turning it into a reliable, safety-rated electrical signal.

Why this connects to robotics: the force/torque and position-sensing portion of Sensata's catalogue is, in principle, directly relevant to robot perception — a robot needs force/torque sensing at its wrists and ankles and position sensing across its joints, and Sensata has decades of automotive-grade experience building exactly those sensor types. The company has also signaled it sees the opportunity: in July 2025 Sensata spent roughly $340 million to acquire a Swiss MEMS fab, securing captive inertial-sensor supply — a vertical-integration move that management framed partly as positioning for emerging robotic-sensor demand.

The honest framing has to lead, though: Sensata is an automotive company. Roughly 70%+ of revenue is automotive, and the investment case is overwhelmingly an auto-and-industrial sensor turnaround story — a cheap stock, a new CEO, a margin-recovery program — with robotics as a thin, early, optional layer on the side. Of the five names in this sensing batch, Sensata has the weakest direct robotics linkage. You are buying a value-and-turnaround situation with a robotics call option attached, not a robotics pure-play, and the frontmatter conviction reflects that.

Product roadmap

Sensata's roadmap is less a sequence of dated flagship launches and more a steady content-expansion program across the vehicle and industrial platforms it already serves. The strategic centerpiece in the recent period is the July 2025 acquisition of a Swiss MEMS fab for roughly $340 million — this gave Sensata captive supply of MEMS inertial sensors (accelerometers and gyroscopes), which matters both for advanced automotive safety content and as a building block for robotic and industrial inertial sensing. Vertical integration of the MEMS supply is the clearest "roadmap" statement the company has made toward the robotics-adjacent opportunity.

On the core franchise, the roadmap is content-per-vehicle growth driven by electrification. Management has been explicit (Q1 2026 call, April 28, 2026) that Sensata sees content per vehicle on electrified powertrains as roughly double that of an internal-combustion vehicle, and that it expects plug-in-hybrid and extended-range EV production to grow 17% in 2026 with a ~12% CAGR through the rest of the decade — Sensata's high-voltage protection, current sensing and pressure-sensing content rides that. The new CEO, Stephan von Schuckmann, has framed the company's forward plan around three pillars — operational excellence, capital allocation, and growth — which is a turnaround roadmap more than a product roadmap. What Sensata does not make is the robot, the actuator, vision systems, or magnetic joint-position sensor ICs (Allegro's lane). Its robotics-relevant assets are its force/position-sensing know-how and its newly captive MEMS inertial supply — real capabilities, but not yet a named, dated robotics product line.

The financial print

Sensata closed full-year 2025 with revenue of $3.70 billion, down 5.8% year-over-year, and adjusted operating income of $704.9 million at a 19.0% margin — flat margin on

lower revenue, reflecting an auto-production-driven top-line decline partly offset by the cost discipline of the turnaround. Full-year adjusted EPS was $3.42. The Q1 2026 print on April 28, 2026 showed the inflection: revenue of $935 million, up 3% year-over-year (organic growth +4.2%), with adjusted operating income of $174 million at an 18.6% margin, GAAP EPS of $0.59 (up 25.5% YoY) and adjusted EPS of $0.86 (up 10.3% YoY). Free cash flow was $104.6 million at 83% conversion. Management guided Q2 2026 to revenue of $950–980 million and adjusted EPS of $0.89–0.95, with adjusted operating margin expanding to 19.2–19.4%.

The story in the numbers is a turnaround gaining traction: revenue has returned to growth, margins are expanding, and free-cash-flow conversion is strong. Forward consensus has 2026 revenue modestly above 2025 with continued margin expansion and EPS growth into the high-$3s to ~$4 range — the active sell-side names include Truist (which upgraded ST to Buy in May 2026 citing an AI/data-center opportunity), Morgan Stanley, and others, clustering around a recovery-plus-margin-expansion view. At $47.88 the stock trades at a forward P/E of just 11.9 with a market cap of $6.96B — by far the cheapest valuation in this batch (VPG 66x, Keyence 41x, Cognex 37x, Allegro 31x), which is the entire value-case argument. The 1-year return has been strong into mid-May 2026, with the stock recently hitting a 52-week high around $46. The next binary is Q2 2026 earnings, expected around July 28, 2026.

Customer mix today

Sensata’s revenue is dominated by automotive — roughly 70%+ of total revenue sits in the Automotive segment, with Industrials at roughly 20% and Aerospace, Defense and Commercial Equipment at roughly 10% (segment proportions on a reporting basis from the company’s three-segment structure). Within automotive, Sensata sells pressure, temperature, force, position and electrical-protection content to tier-1 suppliers and OEMs globally; within industrials it serves HVAC/refrigeration, heavy vehicle and general industrial; the aerospace/defense segment is the smallest but the highest-margin.

The 2024-versus-2026 change is the turnaround narrative more than a mix shift: revenue fell 5.8% in 2025 on weak global auto production, then returned to growth in Q1 2026 (+3% reported, +4.2% organic). The structural mix story management is selling is electrification content — electrified vehicles carry roughly double the Sensata content of an ICE vehicle, and the PHEV/EREV production growth (guided +17% in 2026) is the lever. Robotics, candidly, does not show up as a disclosed customer or segment at all — it is embedded, to whatever small extent it exists today, inside the Industrials segment, and it is best understood as a future option enabled by the MEMS-fab acquisition rather than a current revenue contributor. Of the five names in this batch, Sensata’s

customer mix has the least robotics content and the most automotive-cycle exposure — that concentration is the central risk and the reason the multiple is low.

What’s actually happening in the turnaround and robotics positioning

The mechanism that matters for Sensata in 2026 is the turnaround, not robotics. CEO Stephan von Schuckmann, relatively new in the seat, has organized the company around three pillars — operational excellence, capital allocation, and growth — and the Q1 2026 results gave it “more proof points,” in his words: organization-wide operational discipline, a delivering productivity engine, and accelerating strategic initiatives. The concrete evidence is margin expansion on returning revenue and 83% free-cash-flow conversion. The growth thesis underneath is electrification content-per-vehicle: ~2x content on electrified powertrains, with PHEV/EREV production guided up 17% in 2026 and a ~12% CAGR through the decade.

On robotics specifically, the honest assessment is that Sensata is positioning, not yet participating at scale. The $340 million Swiss MEMS-fab acquisition (July 2025) is the real move — it secures captive inertial-sensor supply, which is a genuine building block for robotic and industrial inertial sensing and a hedge against wafer-supply volatility as the robotic-sensor market expands. Truist’s May 2026 upgrade to Buy was explicitly framed around an “AI opportunity,” which captures the market starting to give Sensata credit for data-center and adjacent-growth optionality. But there is no disclosed robotics design win, no named robotics customer, no robotics revenue line. The robotics linkage is “Sensata has force/position-sensing IP and now-captive MEMS supply, and management says it sees the opportunity” — which is a credible setup, not a demonstrated business. The thing to watch on the FY2026 calls is whether “robotics” or “humanoid” moves from a positioning statement to a disclosed design win or revenue contribution; until it does, the robotics case here is the thinnest in the batch.

The competitive threat

Sensata competes across a fragmented field that varies by product domain. In pressure and force/position sensing for automotive and industrial, named competitors include Bosch, Continental, Infineon, TE Connectivity, Amphenol and Honeywell, plus category specialists; in electrical protection, the competitor set includes Littelfuse and Eaton; competitor screens also list Woodward, INFICON, Keysight, Teradyne and others among

Sensata's broader peer group. There is no single dominant competitor and no active IP litigation to flag — this is a scale-and-relationship business where Sensata's installed automotive content and tier-1 relationships are the moat, not a patent estate.

The competitive risk for the thesis is twofold. First, in the core automotive business, Sensata competes for content on every new vehicle platform, and a loss of sensor sockets to Bosch, Continental or TE on a major platform would be material given the 70%+ auto concentration. Second — and more relevant to the robotics framing — in any future robotics push, Sensata would be a late entrant against companies that are already there: VPG and the dedicated robot-sensing specialists in force/torque and tactile, Allegro and Infineon in joint position sensing. Sensata's force/position-sensing heritage is real, but it has not demonstrated a robotics-specific product or design win the way Allegro has (two robotic-joint design wins) or VPG has (shipping to multiple humanoid developers). The competitive bear case is not that Sensata gets displaced in autos overnight — its content is sticky — but that it arrives at the robotics opportunity behind the names that already have the sockets.

The terminal risk

The terminal risk for Sensata is straightforward and it is cyclical-structural rather than technological: the ~70%+ automotive concentration means a sustained global auto-production downturn, or a disorderly EV-transition mix shift, overwhelms the slow industrial-and-robotics diversification before it can scale. Sensata's 2025 revenue decline of 5.8% was a preview — when global auto build rates weaken, Sensata's top line weakens with them, and no amount of cost discipline fully offsets a volume problem. The EV transition is a double-edged sword: management's thesis is that electrified vehicles carry ~2x Sensata content, which is bullish, but the path of the transition is volatile, and a period where ICE volumes fall faster than PHEV/EV volumes ramp would create an air pocket in content growth.

The good news is that this is not an obsolescence risk — Sensata's pressure, temperature, force and position sensors are not being technologically displaced; they are required content on essentially every vehicle and a great deal of industrial equipment. The terminal risk is therefore a valuation-and-growth risk, not an impairment risk: at 11.9x forward earnings the stock is already priced as a low-growth, cyclically-exposed industrial, so the downside from the terminal risk is "the multiple stays low and the turnaround disappoints" rather than "the franchise is destroyed." The constraint this puts on the upside: for Sensata to re-rate meaningfully above a low-teens multiple, it needs to demonstrate that the industrial, aerospace and — eventually — robotics diversification can reduce the auto-cycle dependence. Until then, the cheap multiple is cheap

for a reason.

Bull / Gap / Optionality (Photoncap framing)

1. The cheapest valuation in the batch by a wide margin. At 11.9x forward earnings, Sensata trades at roughly a third of Cognex's multiple and well under a fifth of VPG's. The Q1 2026 turnaround proof points — revenue back to growth, margin expanding to 18.6%, 83% FCF conversion (company release, April 28, 2026) — are landing against a valuation that prices in almost no improvement. The asymmetry is real.

2. The turnaround is delivering measurable proof points. Q1 2026 adjusted EPS of $0.86 was up 10.3% YoY, GAAP EPS up 25.5%, with Q2 guidance for revenue of $950–980M and margin expansion to 19.2–19.4%. CEO von Schuckmann's three-pillar program (operational excellence, capital allocation, growth) is showing up in the numbers, not just the slides.

3. Electrification content-per-vehicle is a structural top-line lever. Management's thesis that electrified powertrains carry ~2x Sensata content, with PHEV/EREV production guided +17% in 2026 and a ~12% CAGR through the decade (Q1 2026 call), gives the core auto business a secular growth driver that partly offsets the cyclical-volume risk.

4. The MEMS-fab acquisition is a credible robotics/inertial-sensing building block. The July 2025 ~$340M Swiss MEMS-fab purchase secures captive inertial-sensor supply — a genuine asset for robotic and industrial inertial sensing and a hedge against wafer volatility. It is the clearest sign management is positioning for the robotic-sensor expansion.

5. Sell-side is starting to give it credit. Truist upgraded ST to Buy in May 2026, explicitly citing an AI/data-center opportunity — an early sign the market is willing to look past the pure auto-cyclical framing toward the diversification and adjacent-growth optionality.

Gap

1. ~70%+ automotive concentration is the dominant risk. Sensata's 2025 revenue fell 5.8% on weak global auto production — when build rates weaken, the whole company weakens. The robotics and industrial diversification is far too small to offset an auto downturn, and the cycle, not the theme, drives the stock near-term.

2. The robotics linkage is the thinnest in the batch. There is no disclosed robotics design win, no named robotics customer, no robotics revenue line — just force/position-

sensing heritage and a MEMS fab. Sensata is positioning, not participating. An investor buying ST for the robotics theme is buying an option that has not yet started paying.

3. The chart is extended on a value name. RSI 70.3 and +24.8% above the 50-day average — Sensata recently hit a 52-week high near $46. The stock has already re-rated some of the turnaround optimism; the cheap multiple is less of a cushion after a 25% run above trend than it looks on the headline P/E.

4. The EV transition is volatile and could create a content air pocket. Management's ~2x-content thesis is sound long-term, but a period where ICE volumes fall faster than electrified volumes ramp would stall content growth. The transition path is not smooth, and Sensata's content tailwind depends on its timing.

Optionality

Event	Date / window	Direction
Q2 2026 earnings vs. $950–980M revenue / $0.89–0.95 EPS guide	~Jul 28, 2026	Binary on the turnaround trajectory
Margin expansion sustains toward 19.2–19.4%	FY2026 quarters	Bull — re-rate catalyst
First disclosed robotics design win or revenue line	2026–2027	Bull — converts the option
Global auto-production build rates	Ongoing	Bear — dominates given 70%+ mix
Further sell-side upgrades on AI/diversification narrative	H2 2026	Bull — multiple-expansion path

---

The trade

Sensata is the value-and-turnaround name of the sensing batch, not the robotics bet — and it has to be sized and framed that way. The structural case: a cheap (11.9x forward) auto-and-industrial sensor company with a delivering turnaround, an electrification content tailwind, and a robotics call option that the MEMS-fab acquisition makes credible

but that has not yet started paying. The two-layer read: layer one, is the thesis real? Yes — but it is a turnaround thesis with a robotics option attached, not a robotics thesis; the conviction reflects the thin theme linkage. Layer two, is the entry acceptable? Moderately — RSI 70.3 and +24.8% above the 50-day average is extended for a value name that just hit a 52-week high, so this is not a back-up-the-truck entry. The trade is to initiate in a zone of current ±5%, roughly $45.49–$50.27, with a clear preference to accumulate on a pullback toward the lower end rather than chase the high; size at 0.75–1.0% of risk capital — a measured position justified by the cheap multiple and the turnaround proof points, but capped because the robotics linkage is the weakest in the batch and the auto-cycle risk is the highest. Stop at roughly $40.00, below the structural base and the rising 50-day cloud; a break there would say the turnaround optimism has unwound. The named catalyst is Q2 2026 earnings around July 28, 2026, where the $950–980M revenue and $0.89–0.95 EPS guide and the margin-expansion path get tested. If you want the force/position-sensing robotics exposure with an actual robotics order book rather than an option, VPG (this batch) has the real humanoid linkage — though at a vastly higher multiple and a far more dangerous chart; if you want it cheap and disciplined, Allegro is the better-positioned value-meets-robotics name. Sensata is the deep-value turnaround you own for the turnaround, with robotics as the free upside. **Conviction: 5 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

VPG — Vishay Precision Group, Inc. · WATCH (Tier-2) · Conv 5/10 · Bucket C

Vishay Precision Group, Inc. (VPG)

A sleepy precision-sensor company that the market has suddenly decided is a humanoid-robotics pure-play — the story is real, the chart is not investable here.

Investment Research · Photoncap-style deep dive · v1 of "VPG" · May 14, 2026

What Vishay Precision Group physically does

Vishay Precision Group makes the components that let a machine know how hard it is pushing, how much it is carrying, and where its joints are. The core technology is the foil strain gauge — a microscopically thin metal pattern bonded to a flexure that changes electrical resistance by a few parts per million when the flexure deforms under load. VPG has made these for decades through its Micro-Measurements brand, and it converts them into finished sensing products: load cells (the Tedea-Huntleigh, Sensortronics, Revere and Celtron brands that go into industrial scales and weigh-bridges), precision resistors (the Foil Resistors business, spun out of the original Vishay heritage), and increasingly force/torque and tactile sensing modules aimed at automation. The company reports in three segments — Sensors, Weighing Solutions and Measurement Systems — and the strain-gauge die sits underneath essentially all of it.

Why this matters for robotics: a humanoid robot has two perception problems. It needs to see (machine vision, covered elsewhere in this batch), and it needs to feel. Feeling is the harder, less-solved problem. To pick up an egg without crushing it, to insert a connector, to walk over uneven ground without falling, a robot needs force/torque sensing at the wrist and ankle and tactile sensing at the fingertips. A six-axis force/torque sensor at the wrist tells the controller the full vector of contact force; a tactile array on the fingertip gives a pressure map across the contact patch. VPG's strain-gauge expertise maps directly onto both — the company has publicly described its push into "fingertip tactile sensing" and a broader "strategic shift to robotics" (Sidoti conference, early 2026).

The honest framing is that VPG is not yet a robotics company. It is a ~$307M-revenue

precision-sensing company with a small, fast-growing robotics order book bolted onto its Sensors segment. The strain-gauge IP is genuinely differentiated — accuracy, drift, hysteresis and temperature stability are hard-won and VPG is among the best in the world at it — but the question for the trade is whether a few million dollars of humanoid orders justifies a stock that has nearly doubled off its 50-day average.

Product roadmap

VPG's robotics push runs through the Sensors segment rather than a single named product line, which makes the roadmap less crisp than a semiconductor-equipment name — but the company has been specific in its 2026 disclosures. The foundation is the established force-sensor and load-cell portfolio (the VPG Force Sensors brands — Tedea-Huntleigh, Sensortronics, Revere, Celtron) plus the Micro-Measurements strain-gauge catalogue, both of which predate the robotics story by decades and continue to serve industrial weighing, avionics and test-and-measurement.

The robotics-specific work is newer. Through 2025 VPG began shipping custom force/torque and tactile sensing modules to humanoid developers, and on the Q1 FY2026 call (May 12, 2026) management quantified the pipeline concretely: VPG has shipped to its first two humanoid customers, taken an initial prototype order from a third, and is in early discussions with a fourth — described as a startup in defense and industrial applications. The second customer placed an order of roughly $1.5M in late 2025, primarily for tactile sensing. Q4 FY2025 humanoid-related bookings were about $0.8M. The company says it has already shipped sensors covering "a few dozen robots," each robot using roughly 20–30 sensors, and expects a follow-on order for hundreds of units in Q2 FY2026, with the potential to ramp to "hundreds of robots per week" by the end of 2026 — that last figure is management aspiration, not booked backlog, and should be read as such.

What VPG does not make is the robot, the actuator, the vision stack or the controller. It is a component supplier, and even within sensing it is a tactile/force specialist — it does not make the magnetic joint-position sensors that Allegro and others supply. The financial model management has put forward assumes a 50% annual revenue growth rate for the humanoid business off a 2025 baseline, which gets the robotics line to "low tens of millions" — confirmed as a management framework, not consensus.

The financial print

VPG closed FY2025 (year ended December 31, 2025) with revenue of $307.2 million, essentially flat at +0.2% year-over-year, and an operating margin of just 4.5%, down from 5.5% in 2024 — results released February 11, 2026. Full-year net earnings were $5.3 million, or $0.40 per diluted share, down from $0.74 in 2024; adjusted EBITDA was $28.2 million at a 9.2% margin; gross margin slipped to 38.9% from 41.0%. This is the uncomfortable part of the story: the core business spent 2025 going sideways at thin and falling margins, with Q4 FY2025 operating margin compressing to 2.3% on unfavorable mix, inventory reductions and higher fixed and logistics costs.

The Q1 FY2026 print on May 12, 2026 is what changed the narrative. Revenue was $84.4 million, up 17.6% year-over-year and roughly 8.3% above the $77.9M consensus, with Sensors segment revenue of $33.3 million up 23.1%. The headline number, though, was bookings: $102.1 million, up 25.5% sequentially, the third-highest quarterly bookings in company history, lifting book-to-bill to 1.21 and backlog to about $125 million. Sensors booked at a 1.36 book-to-bill. Adjusted EPS was still only $0.07 — the orders are running well ahead of the P&L, which is the bull case and the valuation problem in one sentence. Forward consensus for FY2026 revenue sits around $326–329 million per compiled sell-side estimates, up from ~$320M ninety days ago. At $102.94 the stock trades on a forward P/E of 65.8 with a market cap of $1.37B — a multiple that only makes sense if you underwrite the humanoid ramp aggressively. The 1-year return into mid-May 2026 is well over +150%, with the stock up roughly 57% on the Q1 print alone. The next binary is Q2 FY2026 earnings, expected around August 11, 2026.

Customer mix today

VPG’s revenue splits by segment rather than by named customer, and the company does not disclose customer concentration the way a single-product supplier would. On a reporting basis, Sensors is roughly 40% of revenue, Weighing Solutions roughly 35%, and Measurement Systems roughly 25% — these are approximate proportions derived from segment disclosure, not audited customer percentages. The end-market spread inside that is wide: semiconductor equipment, data centers, avionics, military and space, industrial automation, precision weighing and test-and-measurement. Q1 FY2026’s bookings strength was attributed specifically to semiconductor equipment, data centers, avionics, military/space and select industrial markets — robotics was a contributor but not the largest one.

The structural shift the bulls are paying for is the humanoid line going from essentially zero in 2024 to a low-single-digit-million run-rate in 2025 to management’s “low tens of millions” framework for 2026. Put in proportion: even at the optimistic end, humanoid

robotics is plausibly 5–10% of total revenue in 2026 — real, growing fast, but not yet the company. The 2024-versus-2026 change that matters is less the customer mix than the orders mix: a flat-revenue, thin-margin industrial company is suddenly booking at 1.21x with a humanoid pipeline of four named developers. Whether that converts to durable revenue at acceptable margin is the entire question, and Q1 FY2026 did not yet answer it on the P&L.

What’s actually happening at the humanoid developers

VPG has not named its humanoid customers, so this section is necessarily built from the company’s own characterizations on the Q1 and Q4 calls rather than from customer-side confirmation. The picture management painted: customer one and customer two are both past prototype and into follow-on ordering — customer two’s ~$1.5M late-2025 order was “primarily for tactile sensing,” which tells you VPG is winning the fingertip-sensing socket, not just the wrist force/torque socket. Customer three placed an initial prototype order. Customer four is in early discussions and is described as a defense/industrial robotics startup, which is a slightly different profile than a consumer-humanoid OEM.

The mechanism that matters is sensor count per robot. Management’s “20–30 sensors per robot” figure is the leverage: if the humanoid market moves from prototype fleets of dozens to early-production fleets of hundreds, VPG’s unit volume scales 10–20x without winning a single new customer. The Q2 FY2026 “follow-on order for hundreds of units” that management flagged is the near-term proof point — if it lands, the 50%-CAGR framework gets more credible; if it slips, the multiple is indefensible. The thing to watch with genuine skepticism is that every humanoid OEM is simultaneously evaluating whether to build tactile sensing in-house, and several (Tesla Optimus most prominently) have shown a strong vertical-integration instinct. VPG winning prototype sockets is not the same as VPG holding production sockets at scale.

The competitive threat

VPG’s competitive set splits by where you look. In its legacy weighing and force-measurement business, MarketBeat’s competitor screen lists Keysight, Teledyne, Allient, Diversified Technical Systems and Teradyne among ~17 named competitors — a fragmented, mature field where VPG competes on accuracy and brand. In the robotics-specific force/torque and tactile niche, the more relevant competitors are the

dedicated robot-sensing specialists: ATI Industrial Automation (owned by Novanta), Bota Systems, OnRobot, and a wave of Chinese tactile-sensing startups, plus the humanoid OEMs' own internal efforts.

There is no active IP litigation to flag. The competitive risk here is not a courtroom; it is two slower-moving forces. First, commoditization — fingertip tactile sensing is an area of intense academic and startup activity, and a foil-strain-gauge approach is not the only architecture (capacitive, optical/vision-based and MEMS tactile arrays all have credible proponents). VPG's edge is manufacturing maturity and sensor-grade accuracy, but a "good enough and cheap" Chinese array could undercut it at the volumes that matter. Second, the OEMs themselves: a humanoid developer that decides tactile sensing is core IP can design VPG out. The bull rebuttal is that VPG's strain-gauge process is genuinely hard to replicate and that OEMs racing to production would rather buy than build — but that is an argument, not a moat you can measure yet.

The terminal risk

The terminal risk for VPG is that the humanoid tactile-sensing socket it is winning today is a prototype-phase socket, not a production-phase one. The structural transition that would obsolete the current VPG approach is twofold. The first path is vertical integration: if the humanoid winners — and this is a winner-take-most market — bring tactile and force sensing in-house the way Tesla has signaled across much of its actuator and sensor stack, VPG's design wins evaporate at exactly the moment volume arrives. The second path is architectural: foil strain gauges are a mature, ~80-year-old technology, and the next generation of tactile sensing may be vision-based (a camera looking at a deformable gel fingertip, the GelSight/MIT lineage) or MEMS-based. If the industry standardizes on a non-strain-gauge architecture for fingertips, VPG's incumbency in industrial load cells does not transfer.

VPG does have a credible near-term roadmap — it is shipping, iterating with four named developers, and its accuracy advantage is real. But the multiple the market is now paying (66x forward earnings) implicitly assumes VPG both scales with the humanoid market and holds its sockets through the production transition. Neither is proven, and the terminal risk is precisely that the company is a prototype-era supplier being priced as a production-era one. This is the constraint on how much you can pay — and at current levels you are paying a lot.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. The order book has genuinely inflected. Q1 FY2026 bookings of $102.1M, up 25.5% sequentially, were the third-highest in company history, with book-to-bill at 1.21 and backlog at ~$125M — and Sensors specifically booked at 1.36 (company release, May 12, 2026). This is not a robotics-only story; it is broad strength across semi equipment, data centers and avionics. Even setting humanoid aside, the core business has visibly turned, which de-risks the floor.

2. The humanoid pipeline is real, named and widening. VPG has shipped to two humanoid customers, has a prototype order from a third, and is in discussions with a fourth (Q1 FY2026 call). Customer two's ~$1.5M order was "primarily for tactile sensing" — VPG is winning the hardest, highest-value sensing socket, not the commodity one. Four developers in the pipeline at this early stage is a credible spread of bets.

3. Sensor-count-per-robot is the operating leverage. At management's stated 20–30 sensors per robot, VPG's unit volume scales 10–20x as the humanoid market moves from prototype dozens to production hundreds — with no new customer wins required. Management's 50%-CAGR humanoid framework off a 2025 baseline (Q1 FY2026 call) gets the line to "low tens of millions"; the Q2 "hundreds of units" follow-on order is the near-term test.

4. Operating leverage on the legacy business is dormant but real. FY2025 operating margin was a thin 4.5%, depressed by mix and inventory actions — but VPG ran 41% gross margin as recently as 2024. A bookings-driven revenue recovery against a fixed cost base should expand margins meaningfully through 2026, which is why FY2026 consensus revenue of ~$326–329M matters more than the trailing print.

5. Cross-theme relevance to Photonics (Theme 3). VPG's precision-resistor and strain-gauge content also feeds optical and photonics instrumentation, giving it a second secular tailwind beyond robotics — a modest diversification of the growth story that the pure-play framing understates.

Gap

1. The valuation has fully detached from the P&L. Q1 FY2026 adjusted EPS was $0.07 — unchanged year-over-year — while the stock trades at 65.8x forward earnings and jumped ~57% on the print. The orders are running far ahead of earnings, and the market is paying a production-scale multiple for a company still earning prototype-scale profits. Any wobble in the ramp story re-rates this hard.

2. The chart is one of the most overbought in the entire robotics universe. RSI

90.5 and +94.6% above the 50-day moving average is a statistically extreme reading — the stock has nearly doubled relative to its own recent trend. Even in a secular theme, mean reversion from a 90+ RSI is the base case; the question is not if but how deep.

3. The core business margin structure is weak. FY2025 operating margin of 4.5% and Q4 FY2025 of 2.3% show a company with little cushion. Humanoid revenue at "low tens of millions" is not yet enough to move consolidated margin, and if industrial/semi bookings normalize before humanoid scales, VPG is back to being a thin-margin industrial at a rich multiple.

4. Customer concentration and disclosure opacity. VPG does not name its humanoid customers or disclose their order cadence in detail, so the bull case rests heavily on management characterization. With a winner-take-most humanoid market, losing even one of the two shipping customers — to vertical integration or to a competitor — would be material and the market would learn about it late.

Optionality

Event	Date / window	Direction
Q2 FY2026 earnings — humanoid "hundreds of units" follow-on order	~Aug 11, 2026	Binary on the ramp thesis
Fourth humanoid customer converts from discussion to prototype order	H2 2026	Bull if delivers
Core Sensors/semi-equipment bookings sustain 1.2x+ book-to-bill	Q2–Q3 FY2026	Bull — de-risks the floor
A named humanoid OEM announces in-house tactile sensing	Anytime	Bear — terminal-risk confirmation
Margin recovery on FY2026 revenue toward ~$326–329M consensus	FY2026 reporting	Bull if OPM expands past high-single-digits

---

The trade

VPG is the rare case where the thesis and the timing point in opposite directions, and price discipline has to win. The two-layer framework: Layer one — is the structural thesis real? Yes, qualifiedly. VPG has a genuine, widening humanoid tactile-sensing order book, the sensor-count-per-robot leverage is real, and the core business has visibly inflected on bookings. This is an early-cycle robotics-sensing name worth owning at some price. Layer two — is the entry timing acceptable? No. RSI 90.5 and +94.6% versus the 50-day average is among the most extended readings in the entire robotics universe, on a stock that just gapped ~57% on a print where EPS did not move. Paying 66x forward earnings into a 90+ RSI is buying the narrative at the moment of maximum narrative. The discipline is to put VPG on the sheet as a Bucket C name — significantly extended, chase risk explicit — and wait. The technically-honest entry zone is current ±5%, i.e. roughly $97.79–$108.09, but that should be read as where you would be a forced buyer only after a meaningful RSI reset, not where you initiate today; a more realistic accumulation zone is the post-gap consolidation back toward the low-to-mid $80s if the tape gives it. Size small — 0.25–0.50% of risk capital — precisely because the entry is poor and the position is a toe-hold on a structural theme, not a high-conviction initiation. Stop at roughly $78, below the pre-gap structural base; a break there says the bookings inflection was a one-quarter event. The named catalyst is Q2 FY2026 earnings around August 11, 2026, where the “hundreds of units” humanoid follow-on order either confirms the ramp or doesn’t. If you want a cleaner expression of the robotics-sensing thesis without VPG’s valuation and chart extension, Allegro (ALGM) — also in this batch — offers humanoid joint-sensing exposure at 31x forward earnings on a far less stretched chart. **Conviction: 5 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

ARBE — Arbe Robotics Ltd. · SKIP / WAIT (Tier-3) · Conv 3/10 · Bucket C

Arbe Robotics Ltd. (ARBE)

A pre-revenue 4D imaging-radar chip company with a genuine technology lead — and a 2026 production target it has promised, and missed, before.

Investment Research · Photoncap-style deep dive · v1 of "Arbe Robotics" · May 14, 2026

What Arbe physically does

Arbe designs the silicon and software for 4D imaging radar. Conventional automotive radar — the kind already in millions of cars for adaptive cruise control and blind-spot warning — produces a coarse picture: it tells you something is out there and roughly how fast it is closing, but it cannot resolve fine angular detail, so it cannot reliably tell a stationary car from an overhead sign from a manhole cover. That ambiguity is why legacy radar gets "filtered" — software throws away static returns to avoid false braking — and why radar alone has not been trusted for higher levels of autonomy. Arbe's pitch is a radar with enough channels and processing to produce a genuinely high-resolution 4D point cloud: range, azimuth, elevation, and Doppler velocity, at a resolution Arbe claims is an order of magnitude higher than competing radar. "4D" here means the same thing it does in lidar — three spatial dimensions plus per-point velocity — but radar gets it through a different physics, and radar has the structural advantages of working in rain, fog, snow and darkness, and of being far cheaper than lidar.

The product is a chipset, not a finished sensor: Arbe sells a transmitter chip, a receiver chip, and a radar processor chip, plus the perception software that turns the raw returns into a usable point cloud. Tier-1 suppliers and OEMs integrate that chipset into radar modules. Arbe's two reference module designs are Phoenix — a forward-facing radar for mapping and tracking what is ahead — and Lynx, a surround/corner radar for the rear and side positions, which together are pitched as a 360-degree imaging-radar suite. The binding-constraint argument for Arbe is that L2+ and L3 autonomy need a sensor that delivers lidar-like environmental understanding at radar-like cost and radar-like all-weather reliability — and that high-resolution imaging radar is that sensor. The

expansion thesis that puts Arbe in the Robotics theme: the same all-weather, low-cost, velocity-aware perception is useful for off-highway machines, defense platforms and robots, and Arbe has begun explicitly marketing into those adjacencies — including an HD imaging radar variant for off-highway applications.

Product roadmap

Arbe's product nomenclature centers on the chipset and the two module reference designs. The chipset — transmitter, receiver and processor chips — is the actual product Arbe sells; Phoenix (forward-facing imaging radar) and Lynx (surround/corner imaging radar, introduced in 2022 as an "industry-first surround imaging radar") are the reference modules built around it. The critical dated event: Arbe plans the transition of its radar chip from development to production for sale to customers during 2026 — a target the company itself hedges, and one it has set and missed in prior years, which is the single most important piece of context in the entire file. Per Arbe's 2025 annual report (20-F), the company "can give no assurance that this timetable will be met," and one industry account puts Lynx and Phoenix module production at customers beginning as late as late 2027. Arbe has also recently introduced an HD imaging radar variant aimed at off-highway and industrial applications, the clearest product evidence of the move beyond passenger-car automotive into the robotics/heavy-machinery adjacency.

What Arbe does not make: it does not make lidar, cameras or the host compute, and it does not make finished radar modules at scale — it is a chipset-and-software licensor that depends on Tier-1 integrators to build and sell the modules. That dependency is structural: Arbe cannot reach production revenue on its own timeline; it reaches it only when a Tier-1 and an OEM commit a platform.

The financial print

Arbe is, for practical purposes, a pre-revenue company. Full-year 2025 revenue was approximately $1.0 million, with Q4 2025 revenue of about $0.5 million, and the 2025 net loss was $45.2 million on roughly $34.8 million of R&D spend. Management's guidance for 2026 is revenue of $4–6 million and an adjusted EBITDA loss of $28–31 million — so even the bull-case 2026 is a company with mid-single-digit-million revenue burning roughly $30 million. The balance sheet is the existential issue: Arbe held about $45 million in cash as of December 31, 2025, and in January 2026 raised $18.5 million of gross proceeds in an underwritten public offering, while also cutting expenses by roughly 15%.

Combining the year-end cash, the January raise and the guided ~$30 million burn, the runway is on the order of roughly two years if everything goes to plan — but “to plan” includes the 2026 production transition actually happening, and any slip both extends the burn and damages the financing terms of the next raise. There is no forward P/E and the conventional valuation framework does not apply; the roughly $145 million market cap is a venture-style bet on optionality, not a multiple on a business. Q1 2026 earnings are scheduled for May 28, 2026 — the next binary, and the first read on whether the 2026 production transition is on track.

Customer mix today

There is no production customer mix to break out, because there is no production. The honest description of Arbe’s “customer mix” is a pipeline of engagements at varying maturity. Per Arbe’s disclosures, in 2024 the company engaged with roughly 15 OEMs, of which about 11 had advanced to a bid stage and about 8 had entered an “advanced perception project” phase, and Arbe’s radar chipset was selected by a top-10 global automaker to support development of that OEM’s next-generation imaging-radar systems. Those are real engagements, but the language — “bid stage,” “advanced perception project,” “support development” — is the language of pre-production qualification, not signed production contracts with unit volumes. The structural shift Arbe wants investors to see is from a 2024 automotive-only pipeline toward a 2026 pipeline that adds defense and robotics/off-highway adjacencies, evidenced by the new HD imaging radar for off-highway applications. But the financial reality is unambiguous: $1.0 million of 2025 revenue means no customer is yet a meaningful revenue contributor, and the entire investment case is that the pipeline converts.

What’s actually happening at the Tier-1 / OEM conversion point

The mechanism that matters for Arbe is the conversion of “bid stage” and “advanced perception project” engagements into committed Tier-1-plus-OEM production platforms — and the honest read is that this has not happened on the timeline Arbe has repeatedly projected. Arbe’s chipset was selected by a top-10 automaker for next-generation imaging-radar development, and the company says Lynx and Phoenix modules are being integrated into multiple platforms, but the dates keep moving: the chipset production transition is now a 2026 target (hedged), and module production at customers is described in some accounts as beginning as late as 2027. The structural problem is that

Arbe does not control its own revenue timeline — it depends on a Tier-1 integrator building the module and an OEM committing a vehicle platform, and OEM platform decisions are slow, conservative and easily delayed. The bull mechanism, if it works: once a single major OEM platform commits to Arbe's chipset, the per-vehicle content and the multi-year platform life create a real, durable revenue base, and the company has spent years and roughly $35 million a year of R&D building a genuine resolution lead to win exactly that decision. The bear mechanism: every quarter without a signed production platform burns cash, and the longer it takes, the more time Tier-1 silicon incumbents have to close the resolution gap.

The competitive threat / TI, NXP, Infineon and Mobileye

Arbe's competitive problem is that it is a sub-$150 million-market-cap company trying to win a market against some of the best-capitalized semiconductor companies on earth. Per industry profiles, Arbe's main rivals include Texas Instruments, NXP Semiconductors, Infineon and Bosch on the radar-silicon side, plus Mobileye, which is pushing its own imaging-radar chipset (positioned to reach OEMs via Tier-1s such as Valeo and Continental). Continental's ARS540 is already one of the first 4D radars in mass production — seen on the Hyundai Ioniq 5 — representing the "safe," already-qualified choice for a Tier-1 integrator even if it trails Arbe in raw point-cloud density. The competitive dynamic is the classic micro-cap-versus-incumbent trap: Arbe may genuinely have the best resolution, but a Tier-1 choosing a radar chip weighs supplier longevity, balance-sheet durability, automotive-grade track record and roadmap support — and on every one of those dimensions TI, NXP, Infineon and Mobileye outscore a cash-strapped micro-cap. Arbe's "two-horse race against Mobileye for L3 imaging radar" framing is the optimistic version; the pessimistic version is that the incumbents close the resolution gap with engineering budgets Arbe cannot match, and Arbe's technology lead expires before its pipeline converts.

The terminal risk

The terminal risk for Arbe is straightforward and severe: it runs out of money before the technology lead converts to production revenue, or the incumbents erase the lead before conversion. Arbe is burning roughly $30 million a year against a cash position — even after the January 2026 raise — of well under $100 million. The bull case requires the 2026 production transition to happen roughly on schedule and a major OEM platform to commit; if instead the transition slips into 2027 (as some accounts already

suggest for the modules) and the OEM decisions stay in “bid stage,” Arbe faces a dilutive raise from a position of weakness, and the technology window narrows as TI, NXP, Infineon and Mobileye iterate their own 4D radar silicon. The named alternative beneficiaries are precisely those incumbents, plus Continental on the module side. Arbe has a credible roadmap and a real resolution advantage — that is not in dispute — but a credible roadmap is not the same as a financed one, and that is what caps any reasonable position size in this name to a token, venture-style allocation.

Bull / Gap / Optionality (Photoncap framing)

1. The technology lead is real and well-documented. Arbe claims, and industry coverage broadly accepts, an imaging-radar resolution roughly an order of magnitude above competing radar, built over years and roughly $35 million a year of R&D. High-resolution 4D radar at radar cost and all-weather reliability is genuinely the sensor L2+/L3 autonomy wants, and Arbe is the most prominent pure-play.

2. The pipeline exists and includes a top-10 automaker. Arbe disclosed engagement with ~15 OEMs, ~11 at bid stage and ~8 in advanced perception projects, plus a top-10 global automaker that selected its chipset for next-generation imaging-radar development. The funnel is real even if conversion has lagged.

3. The robotics/defense expansion widens the addressable market. The new HD imaging radar for off-highway applications, plus explicit defense and robotics positioning, gives Arbe shots beyond the slow passenger-car platform cycle — all-weather, low-cost, velocity-aware perception is genuinely useful for heavy machinery and robots.

4. The company has bought itself time. The January 2026 raise of $18.5 million plus a roughly 15% expense cut extended the runway to something like two years on the guided burn — enough, if the 2026 production transition lands, to reach the conversion inflection without an immediate emergency raise.

5. The micro-cap structure means a single OEM win re-rates the stock. At a ~$145 million market cap, one signed major-OEM production platform — multi-year, per-vehicle content — would transform the revenue outlook and the equity story. The asymmetry is real; it is just low-probability.

Gap

1. It is effectively pre-revenue with a history of missed production dates. $1.0 million of 2025 revenue, and a 2026 production transition the company itself says it “can

give no assurance" of meeting — a target it has set and missed before. The single most important fact in the file is that Arbe's timelines have not been reliable.

2. The runway is finite and the next raise will likely be dilutive. Roughly $30 million annual burn against well under $100 million of cash even post-raise. If the production transition slips, Arbe raises again from weakness, diluting an already micro-cap share base.

3. The competitors are vastly better capitalized. TI, NXP, Infineon, Bosch and Mobileye all build 4D radar silicon and all outscore Arbe on supplier longevity and balance-sheet durability — the exact criteria a Tier-1 weighs. Continental's ARS540 is already in mass production as the "safe" choice.

4. The robotics revenue is entirely aspirational. The off-highway/defense/robotics positioning is roadmap and marketing — there is no disclosed robotics revenue. Within the Robotics theme, ARBE is the purest example of "option value on a loss-making balance sheet" with the option deepest out of the money.

Optionality

Event	Date / window	Direction
Q1 2026 earnings	May 28, 2026	Binary on 2026 production-transition progress
Radar chipset enters production for sale	During 2026 (hedged)	Bull if on schedule, Bear if slips
First signed major-OEM production platform	2026–2027	Bull — would re-rate the stock
Defense / off-highway design win	2026 ongoing	Bull
Dilutive capital raise	If transition slips	Bear

The trade

Arbe is a Bucket-C, token-size venture bet — a genuine 4D imaging-radar technology lead attached to a pre-revenue income statement and a runway that depends on a 2026 production target the company has missed before. If taken at all, initiate at $1.12–$1.24 (current price ± ~5%), but note a real price-verification flag: secondary sources showed ARBE closing around $1.01 on May 13, 2026, roughly 14–17% below the $1.18

reference — this is a volatile micro-cap that swings double digits on single news items, so verify the live price before committing capital. Size at no more than roughly 0.4% of risk capital — this is the smallest, most speculative position in the batch, appropriate to a name where the realistic outcome distribution is bimodal — and set a stop at $0.78, accepting that a sub-$1.20 micro-cap can gap through any stop. The named catalyst is Q1 2026 earnings on May 28, 2026, the first read on the production transition; the real binary, though, is whether any quarter in 2026 produces a signed major-OEM production platform rather than another “bid stage” update. If you want imaging-radar exposure without the existential balance-sheet risk, the cleaner expression is to own it through a diversified, profitable name — Mobileye carries an imaging-radar chipset on a fortress balance sheet, and indie Semiconductor (INDI) has radar SoC revenue today. Arbe is the pure lottery ticket; treat the position size accordingly. **Conviction: 3 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

Section 6 — Layer 3 — Brain: compute and memory

The robot runs three compute tiers simultaneously, at three different timescales:

- **Tier A — High-level reasoning ("the cortex"), ~10 Hz.** Decides what to do. Runs a vision-language model on a GPU; foundation models like NVIDIA's GR00T and Google's RT-2 sit here. NVIDIA (NVDA) is the de-facto standard via its Jetson Thor module, Isaac robotics platform and Omniverse simulator — same vertical-integration playbook that produced data-centre dominance. Qualcomm (QCOM) is the credible #2 with its Dragonwing IQ family and named design wins at Figure and NEURA Robotics. Combined NVIDIA-plus-Qualcomm share at this tier is probably 80%+.
- **Tier B — Mid-level perception ("the visual cortex"), 30-60 Hz.** Turns raw camera/lidar feeds into a 3D scene representation. Handled by vision SoCs (Ambarella CVflow), low-power FPGAs (Lattice Semiconductor as the "sensor-fusion glue layer" with explicit robotics partnerships at NVIDIA Halos and TI), or by the same Jetson/Dragonwing module running multiple workloads.
- **Tier C — Motor control ("the spinal reflex"), 1,000-10,000 Hz.** Per-millisecond torque commands to each motor; zero latency tolerance. Home of low-power microcontrollers — Ambiq Micro (AMBQ) is positioning here with its subthreshold-voltage edge-AI SoCs, but the tier is broadly fragmented (NXP, STMicro, Texas Instruments, Infineon).

Memory is the often-forgotten layer. Robots have one peculiar memory requirement that conventional compute does not: they must save state instantly when power is interrupted, so they do not drop the object they are holding. This is the role of MRAM (Magnetoresistive RAM), with Everspin Technologies (MRAM) as the only public pure-play. Industrial automation is already a core revenue segment, and a recent $40M defense subcontract anchors the 2026 print.

The structural read on the layer. Three tiers, three vendor groups — NVIDIA and Ambiq are not direct competitors; they sell into different sockets. This basket holds

Tier-1 BUY positions in NVDA, QCOM and LSCC because they sit at three different chokepoints in the brain layer (Tier-A platform, Tier-A alternative, Tier-B glue layer) and each has its own moat. AMBA (Ambarella) is Tier-2 WATCH because the vision-SoC slot is contested. AMBQ (Ambiq) and MRAM (Everspin) are Tier-3 SKIP / deep-pullback only — both extraordinarily extended on the tape (AMBQ +116% vs 50-DMA, MRAM +174%) — the stories are real but the entry discipline is non-negotiable at these levels.

LSCC — Lattice Semiconductor Corporation · BUY (Tier-1) · Conv 7/10 · Bucket B

Lattice Semiconductor Corporation (LSCC)

The low-power FPGA as the robot's sensor-fusion glue layer — explicit robotics design-ins via the NVIDIA Halos and TI partnerships — and unlike the mega-caps, robotics is a real growth vector, not a rounding error.

Investment Research · Photoncap-style deep dive · v1 of "Lattice Semiconductor" · May 14, 2026

What Lattice physically does

Lattice makes small, low-power field-programmable gate arrays — FPGAs that, unlike Intel's or AMD's data-center FPGAs, are measured in milliwatts and millimeters rather than watts and centimeters. An FPGA is reconfigurable logic: a chip whose internal wiring is programmed after manufacture to do exactly the job a system designer needs. In a robot, that job is the unglamorous-but-essential connective tissue — taking the raw outputs of a dozen heterogeneous sensors (cameras, depth sensors, LiDAR, IMUs, force-torque sensors, encoders) and synchronizing, time-stamping, pre-processing and routing them into the main application processor over a clean, low-latency data pipeline. The robot's "brain" (an NVIDIA Jetson or a Qualcomm Dragonwing) is the expensive SoC that runs the model; the Lattice FPGA is the glue layer that feeds it sensor data fast enough and cheaply enough that the brain is not bottlenecked or burdened.

The product families that do this are the Nexus platform — Lattice's low-power FPGA line built on 28nm fully-depleted silicon-on-insulator, including the CrossLink-NX devices optimized for vision and sensor bridging — and the higher-capacity Avant mid-range family. Around the silicon Lattice has built reference solutions: the Holoscan Sensor Bridge, a Lattice-FPGA-based design that creates synchronized, low-latency sensor-data pipelines, is the centerpiece of its robotics push.

Why does a milliwatt-class FPGA become a binding constraint in a robot? Because sensor fusion is a latency-and-power problem that the main SoC is bad at. Offloading depth processing, sensor synchronization and pre-processing onto a dedicated low-power FPGA frees the application processor's compute and power budget for the model — and in a battery-powered humanoid, every watt and every millisecond of latency matters. Lattice's robotics thesis, and the reason this name is the strongest robotics story in this compute-and-materials batch, is that the FPGA is genuinely the right tool for that job, the company has explicitly pivoted to market it that way, and — crucially — Lattice already has a real, growing AI revenue mix in the print, so the robotics angle is incremental upside on a thesis that is already working rather than a pure call option.

Product roadmap

The roadmap has two FPGA platform families and a robotics-specific reference-design layer. On the silicon side: the Nexus platform is the established low-power line — built on 28nm FD-SOI, it spans the CrossLink-NX vision/sensor-bridge devices and the broader small-FPGA portfolio. Lattice extended it with Nexus 2, unveiled December 2024, the next generation of the small-FPGA platform. The mid-range Avant family moved up in capability with the Avant 30 and Avant 50 devices, also introduced December 2024, targeting edge-optimized and advanced-connectivity applications — the parts that handle higher-bandwidth sensor aggregation. Lattice also extended into security with a new secure-control FPGA family carrying crypto-agility and hardware root-of-trust, relevant to robots that need a trusted boot and update path.

The robotics-specific layer is where the 2026 news cluster sits. The Holoscan Sensor Bridge is Lattice's FPGA-based reference design for synchronized sensor pipelines. In March 2026, at NVIDIA GTC, Lattice joined the NVIDIA Halos ecosystem — the physical-AI functional-safety program — committing to build Halos-certified Holoscan Sensor Bridge designs. In April 2026 Lattice announced a collaboration with Texas Instruments to combine TI's sensing technologies with the Lattice Holoscan Sensor Bridge, giving robotics and industrial developers a flexible hardware foundation for synchronized, low-latency sensor pipelines. And at Embedded World 2026 (February 2026) Lattice and Airy3D showcased a humanoid-and-robotic 3D-vision demo combining Airy3D's DepthIQ depth technology with a CrossLink-NX FPGA — depth perception in a very small, low-power form factor, with depth processing offloaded to the FPGA.

What Lattice does not make is the robot brain — it does not make the application SoC or the foundation model, and it explicitly positions itself as the complement to NVIDIA's and TI's silicon rather than a competitor. It also does not make actuators or sensors. The cadence is roughly a two-year platform refresh (Nexus to Nexus 2, Avant capacity

steps) layered with a steady stream of reference designs and ecosystem partnerships.

The financial print

Lattice reported Q1 2026 on May 4, 2026: revenue of $170.9 million, up 42% year-on-year and 17% sequentially, with EPS up more than 80% year-on-year to $0.41, above the high end of guidance. The growth driver was the Compute & Communications segment, up 86% year-on-year and 15% sequentially to a record, with management guiding that roughly 38% of 2026 revenue is expected to come from servers and roughly 25% from AI customers. The other half of the business — Industrial & Automotive, which contains robotics — is the steadier, more diversified leg.

For full-year context, Lattice's revenue troughed through a 2024-2025 inventory-correction cycle and is now in a clean recovery — the 42% year-on-year Q1 2026 growth is the inflection. Sell-side coverage is broad for a mid-cap; the analysts who cover the small-FPGA space (Needham, Stifel, Susquehanna, KeyBanc among the historical voices) are generally constructive on the AI-attach and edge-AI mix story, with the debate centered on the valuation. Forward P/E is approximately 53.5x — a growth multiple, but notably less stretched than Ambarella's ~75x, and the technicals are the calmest in this batch: RSI 57.6 is mid-range, +16.5% above the 50-day moving average is a healthy uptrend rather than a parabolic extension. That relative restraint matters — LSCC is the one name in this compute-and-materials group you can initiate without fighting an overbought tape.

The binary event is **Q2 2026 earnings, expected early August 2026** (Lattice reports on a roughly quarterly cadence; August 3, 2026 is the working estimate, to be confirmed). The read-through this theme cares about is any incremental disclosure on the Industrial segment and traction on the Holoscan Sensor Bridge robotics design-ins.

Customer mix today

Lattice reports by segment rather than by named customer. The structural shift is the headline: the company has flipped from an industrial-and-consumer-weighted revenue base toward Compute & Communications, which is now the larger and faster-growing segment — up 86% year-on-year in Q1 2026, with roughly 38% of 2026 revenue expected from servers and roughly 25% from AI customers (company guidance, Q1 2026). The other major segment, Industrial & Automotive, is roughly the other half of revenue and is where robotics lives.

The 2024-to-2026 change is the story. In 2024, Lattice was a low-power FPGA company emerging from an inventory correction, with the AI/server attach a promising-but-small contributor. By 2026, AI customers are roughly a quarter of revenue and servers nearly 40% — the AI-attach thesis is no longer speculative, it is in the print. Robotics, sitting inside Industrial & Automotive, is the next leg the company is explicitly building toward, and the cluster of 2026 partnerships — NVIDIA Halos in March, TI in April, the Airy3D humanoid 3D-vision demo in February — is the deliberate effort to convert robotics from a use-case slide into a customer-mix line. Lattice has not yet sized robotics as a revenue percentage, so the honest framing is: the AI/server mix shift is confirmed and in the numbers; the robotics mix shift is in the design-in and partnership stage, expected to show up in the Industrial segment over 2026-2027.

What’s actually happening at the NVIDIA Halos and TI partnerships

The mechanism to watch is reference-design adoption. Lattice does not need to win a robot’s main SoC socket — it needs the Holoscan Sensor Bridge to become a standard building block in robot designs, the way a particular power-management or connectivity chip becomes a default choice. The two 2026 partnerships are precisely the channels for that. Joining the NVIDIA Halos ecosystem (March 2026, at GTC) means Lattice’s Holoscan Sensor Bridge designs get a path to Halos functional-safety certification — and because NVIDIA Jetson is the default robot brain, being the certified sensor-bridge complement to Jetson is a powerful pull-through position: a robot designer building on Jetson now has a Lattice-based, pre-certified sensor pipeline as the path of least resistance.

The TI collaboration (April 2026) is the second channel: combining TI’s sensing front-ends with the Lattice Holoscan Sensor Bridge gives industrial and robotics developers a turnkey synchronized sensor pipeline. TI’s reach into industrial and robotics design is enormous, so this is distribution leverage. Be specific about what this is and is not: these are partnerships and reference designs, not disclosed dollar-sized contracts or named humanoid production wins. Lattice has not said “robot OEM X has designed in CrossLink-NX for Y million units.” What it has is structural positioning — the certified complement to the dominant brain (NVIDIA) and a turnkey design with the dominant industrial-sensing vendor (TI). The qualification-to-revenue conversion runs through 2026-2027 as robot programs that adopt these reference designs move to production. The Airy3D 3D-vision demo is the proof-of-concept that the technical fit is real; the partnerships are the go-to-market; the revenue is the still-to-come part.

The competitive threat / the integrated-SoC squeeze and Microchip/Altera

Lattice's competitive picture has two layers. The direct-competitor layer is the low-power FPGA market itself: Microchip (via its PolarFire and IGLOO low-power FPGA lines) and Altera (the former Intel FPGA business, now independent) are the named competitors for the discrete-FPGA socket. Lattice's defense here is genuine product leadership at the lowest-power, smallest-form-factor end — the Nexus and CrossLink-NX parts are purpose-optimized for exactly the milliwatt-class vision-and-sensor-bridge role a robot needs — and it has historically held strong share at that node. This is a real competitive set but not an existential one; it is a share fight within a category Lattice leads.

The more important layer is the integrated-SoC squeeze. The structural risk is not another FPGA vendor — it is the robot brain itself. If NVIDIA's Jetson and Qualcomm's Dragonwing integrate enough sensor I/O, synchronization and bridging function directly onto the SoC, the discrete low-power FPGA gets designed out of the bill of materials. There is no IP litigation at issue; the threat is roadmap integration. The reason this is a manageable rather than terminal risk — and why Lattice's Halos and TI partnerships are smart — is that Lattice has positioned itself as the complement rather than the competitor: NVIDIA itself blessed the Holoscan Sensor Bridge approach by bringing Lattice into the Halos ecosystem, which is a signal that NVIDIA wants a flexible FPGA glue layer in the design rather than absorbing every sensor function onto Jetson. Lattice's bet is that sensor heterogeneity is permanent — robots will always have a changing mix of sensors that a fixed-function SoC cannot anticipate — and reconfigurable logic is the durable answer to that. That bet is plausible but it is the bet.

The terminal risk

The terminal risk is SoC integration removing the discrete FPGA from the robot. If the robotics-compute industry converges on a small number of highly integrated SoCs that absorb sensor-bridge, synchronization and pre-processing functions on-die, then the low-power FPGA — Lattice's whole franchise in this application — loses its socket. The transition window is the 2027-2030 period as humanoid and AMR bills of materials standardize and SoC vendors push integration to lower system cost. The named beneficiaries of that transition are NVIDIA and Qualcomm, the very SoC vendors Lattice

currently partners with.

Lattice's credible defense is structural: sensor heterogeneity and the need for late-stage design flexibility are real and arguably permanent in robotics — a fixed-function SoC cannot anticipate every sensor combination a robot designer will want, and reconfigurable logic is the natural hedge. NVIDIA's decision to bring Lattice into the Halos ecosystem is evidence that even the dominant SoC vendor wants an FPGA glue layer in the architecture. The honest read: this is a slower-burning and more contestable terminal risk than the socket-absorption threat facing a vision-only vendor like Ambarella, because Lattice's function (flexible glue) is harder to fully absorb than a vision block. But it still constrains the multiple — you cannot pay an unlimited price for a company whose role could be integrated away over a five-to-eight-year horizon.

Bull / Gap / Optionality (Photoncap framing)

1. **The AI/server mix shift is already in the print — robotics is incremental upside, not the whole thesis.** Roughly 25% of 2026 revenue is from AI customers and ~38% from servers (company guidance, Q1 2026), and Compute & Communications grew 86% year-on-year. Unlike the mega-caps where robotics is a 1% rounding error, or the pure-plays where robotics is a pure call option, LSCC has a working AI thesis in the numbers and the robotics angle stacks on top of it.

2. **The robotics positioning is the most concrete in this batch.** The NVIDIA Halos ecosystem membership (March 2026), the TI sensor-pipeline collaboration (April 2026), and the Airy3D humanoid 3D-vision demo (February 2026) are three distinct, dated, named go-to-market channels — Lattice is the certified sensor-bridge complement to the dominant robot brain and has turnkey distribution through the dominant industrial-sensing vendor.

3. **The financial inflection is clean and the technicals are the calmest in the group.** Q1 2026 revenue up 42% year-on-year, EPS up over 80%, above the high end of guidance — a clear recovery inflection out of the 2024-2025 inventory correction. And at RSI 57.6 and +16.5% above the 50-day moving average, LSCC is the one name here you can initiate without chasing a parabolic tape.

4. **Lattice is the complement, not the competitor, to the robot brain.** By positioning the Holoscan Sensor Bridge as the FPGA glue layer that NVIDIA and TI want in the design, Lattice converts the integrated-SoC threat into a partnership — NVIDIA bringing Lattice into Halos is the dominant SoC vendor explicitly endorsing a discrete-FPGA role.

5. **The valuation, while a growth multiple, is the most reasonable of the high-multiple names here.** A forward P/E of ~53.5x is full but materially below Ambarella's ~75x, Everspin's ~84x or MP's ~49x — and it sits against a confirmed 42% growth rate with the robotics optionality not yet priced in.

Gap

1. **Robotics is not yet a sized revenue line.** For all three partnerships, Lattice has not disclosed a dollar-sized robotics contract or a named humanoid production design-in. Robotics sits inside the Industrial & Automotive segment as design-ins and reference designs — the conversion to revenue is a 2026-2027 expectation, not a current number.

2. **The integrated-SoC squeeze is the structural overhang.** If NVIDIA's Jetson and Qualcomm's Dragonwing absorb sensor-bridge function on-die, the discrete low-power FPGA loses its socket over 2027-2030. Lattice's "complement not competitor" positioning is smart, but it depends on SoC vendors choosing not to integrate — a choice they control, not Lattice.

3. **The current growth is server/AI-driven, which is its own cycle risk.** Compute & Communications up 86% year-on-year is great, but it ties a large and growing share of revenue to the AI-server capex cycle — if hyperscaler FPGA-attach spending cools, the headline growth rate decelerates hard, independent of anything happening in robotics.

4. **The FPGA competitive set is real.** Microchip (PolarFire, IGLOO) and the independent Altera both compete for the low-power FPGA socket. Lattice leads at the lowest-power node, but it is a share fight, and pricing pressure in a contested category caps margin expansion.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~August 3, 2026 (to be confirmed)	Binary on AI-attach growth; Industrial/robotics read-through
Halos-certified Holoscan Sensor Bridge designs reaching production	2026-2027	Bull — converts the NVIDIA partnership to revenue
TI collaboration turnkey designs adopted by robot/industrial OEMs	2026-2027	Bull

Event	Date / window	Direction
First named humanoid/AMR design-in disclosure	2026-2027	Bull — sizes the robotics opportunity
SoC-integration roadmap moves by NVIDIA/Qualcomm absorbing sensor-bridge	2027-2030	Bear — socket-absorption risk

---

The trade

LSCC is a Bucket B name and the strongest robotics-thesis name in this compute-and-materials batch: a low-power FPGA franchise with a confirmed AI/server mix shift already in the print, the most concrete robotics go-to-market positioning of the group, and — uniquely here — a calm, non-extended tape. Initiate in a **$117-130 entry zone** (current $123.63 minus roughly 1xATR to plus 5%), size at **1.5-2.5% of risk capital** — the cleaner technicals and the in-the-print AI thesis support a fuller position than the extended names — with a **stop near $104** (below the 50-day moving average and the prior consolidation shelf). The defining near-term binary is **Q2 2026 earnings, expected around August 3, 2026**, where the read-through is the AI-attach growth rate and any Industrial-segment robotics commentary. If you want a higher-beta expression of the same physical-AI sensor thesis you take the SoC names (NVDA, QCOM); if you want the dedicated-vision angle, Ambarella (AMBA) — but LSCC is arguably the cleanest risk/reward in the batch because the base thesis already works and robotics is free optionality on top. LSCC's role in a robotics book is a core position, not a satellite — the one name here where the entry tape, the valuation and the thesis-clarity all line up. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

NVDA — NVIDIA Corporation · BUY (Tier-1) · Conv 7/10 · Bucket B

NVIDIA Corporation (NVDA)

The “robot brain” and the robotics operating system in one company — Jetson Thor compute plus the Isaac/GR00T/Omniverse software stack — even though robotics is still a rounding error on the P&L.

Investment Research · Photoncap-style deep dive · v1 of “NVIDIA” · May 14, 2026

What NVIDIA physically does

Strip away the data-center GPU business that everyone already knows and NVIDIA’s robotics franchise is three distinct things stacked on top of each other. At the bottom is silicon you can bolt into a physical machine: the Jetson family of system-on-modules, which are not data-center cards but small, power-constrained compute bricks that run the perception, planning and control loops of a robot in real time, on-board, at 40-130 watts. The current generation, Jetson AGX Thor, is built on the same Blackwell architecture as the data-center parts and delivers up to 2,070 FP4 TFLOPS — roughly 7.5x the AI compute of the prior Orin generation — which is what lets a humanoid run a multi-modal vision-language-action model locally instead of round-tripping to a server. This is the literal “brain” of the robot.

The middle layer is Isaac, NVIDIA’s robotics development platform: the Isaac ROS perception libraries, Isaac Lab for reinforcement-learning training, Isaac Sim for physics-accurate simulation, and the Holoscan sensor pipeline. The top layer is the model layer — Project GR00T, NVIDIA’s foundation-model effort for humanoid robots, plus Cosmos, a “world model” that generates synthetic training data. The strategic point is that NVIDIA is trying to be the Android of robotics: own the chip, own the dev tools, own the simulator, own the base model, and let everyone else build the actual robot on top. That is why the company describes its addressable opportunity as “physical AI” rather than “robot chips.”

Why does this make NVIDIA a binding constraint of the robotics cycle rather than just a participant? Because the hard problem in humanoid robotics right now is not the

actuators or the chassis — it is getting a model that generalizes across tasks to run fast enough, cheap enough and cool enough to sit inside a battery-powered machine. NVIDIA is the only vendor today that offers a fully integrated answer to that problem from silicon to foundation model, and it has the CUDA software lock-in that made its data-center business unassailable. The honest caveat, developed in the financial section below, is that none of this is yet material to the income statement. Robotics-plus-automotive is reported as a single segment and is roughly 1% of company revenue. The robotics thesis on NVDA is a call option embedded in a stock whose price is set almost entirely by data-center GPUs.

Product roadmap

On the silicon side, the lineage runs Jetson Orin (the workhorse of the 2023-2025 robot generation) to the Blackwell-based Jetson Thor family launched in 2026. Jetson AGX Thor developer kits and T5000 production modules reached general availability in August 2025 — the dev kit at $3,499, the T5000 module at $2,999 in 1,000-unit volume. NVIDIA followed with the lower-cost Jetson T4000 module (1,200 FP4 TFLOPS, 64GB, priced at $1,999 at 1,000 units), announced around GTC, which broadens the family down-market into mid-tier autonomous mobile robots, and IGX Thor, a functional-safety-certified variant for industrial and medical edge AI. The pricing ladder — roughly $2,000 to $3,500 per brain — matters because it sets the silicon content per robot at a level that is meaningful if humanoid unit volumes ever reach the millions.

On the software and model side: Isaac GR00T N1, billed as the first open, customizable foundation model for humanoid reasoning, was announced at GTC March 2025. The cadence since has been rapid — GR00T N1.6, then GR00T N1.7 in early access with commercial licensing as of GTC 2026 (March 2026), the latter trained with 20,000 hours of “EgoScale” human video data for better language-following and dexterous control. NVIDIA has publicly committed to GR00T N2 by end-2026, built on a new in-house “DreamZero” world-action model architecture that the company claims will more than double task success rates in unfamiliar environments. Around these sit Cosmos (the synthetic-data world model) and the Halos safety initiative announced at GTC 2026, an ecosystem program for functional-safety certification of physical-AI systems — Lattice Semiconductor and others have joined it.

What NVIDIA does not make is the robot. It does not build actuators, harmonic drives, force-torque sensors, batteries or chassis, and it does not make the rare-earth magnets that go into the motors. It is a compute-and-software layer that needs a thriving ecosystem of robot OEMs to monetize — which is both the bull case (platform economics) and a dependency.

The financial print

NVIDIA closed fiscal 2026 (year ended late January 2026) with revenue of approximately $215.9 billion, with data-center networking up 142% and the data-center segment the overwhelming driver. The most recent quarterly print, Q4 FY26 reported February 25, 2026, delivered roughly $68 billion in revenue. The prior Q1 FY26 print (reported May 2025) showed $44.1 billion in revenue, up 69% year-on-year, with data center at $39 billion — but that quarter carried a $4.5 billion charge tied to H20 export-control inventory write-downs and another $2.5 billion of H20 revenue that could not ship, a reminder that China policy is a live swing factor.

Guidance for Q1 FY27 is revenue of $78.0 billion plus or minus 2%. Wall Street consensus sits modestly above the midpoint — unusual, since NVIDIA typically guides conservatively — with the Street looking for roughly $1.77 in EPS and about 78% year-on-year revenue growth; Goldman Sachs is the visible bull at roughly $80.05 billion, about $2 billion above consensus. The consensus rating is Strong Buy with an average 12-month target around $272, implying roughly 17-22% upside from current levels. Forward P/E is approximately 20.5x — strikingly undemanding for a company growing this fast, which tells you the market is pricing meaningful deceleration risk.

The binary event of the quarter is **Q1 FY27 earnings on May 20, 2026, after the close.** Robotics will not move the print — but management commentary on the robotics/physical-AI roadmap, Jetson Thor traction and any disclosed design-win count is the read-through this theme cares about.

Customer mix today

NVIDIA does not break out robotics customers as revenue percentages because they are not yet material — so the honest framing is design wins, not dollars. On the data-center side that sets the stock price, the hyperscalers (Microsoft, Amazon, Google, Meta, Oracle and the neoclouds) collectively represent well over half of data-center revenue on a reporting basis, with concentration high enough that NVIDIA discloses customer-concentration risk in its filings.

On the robotics side, the named early adopters of the Jetson/Isaac stack as of the August 2025 Jetson Thor GA announcement include Agility Robotics, Amazon Robotics, Boston Dynamics, Caterpillar, Figure, Hexagon, Medtronic and Meta. NVIDIA also disclosed that over 2 million developers now use its robotics stack. The structural shift worth

highlighting: in 2023-2024 the robotics conversation was about Orin in drones and industrial AMRs; by 2026 it is about Thor in humanoids, and the customer list has moved from niche robotics specialists toward industrial heavyweights (Caterpillar partnered with NVIDIA on physical AI for heavy industry) and the best-funded humanoid startups (Figure). The automotive-plus-robotics segment posted roughly $567 million in Q1 FY26 and $586 million in Q2 FY26, up 69% year-on-year — real growth, but on a base that is about 1% of company revenue per multiple sell-side estimates.

What's actually happening at the humanoid OEMs

The mechanism to watch is platform capture: NVIDIA wants every credible humanoid program to standardize on Jetson for compute and Isaac/GR00T for the software stack, the same way the data-center business standardized on CUDA. The evidence that this is happening is the breadth of the GTC 2026 robot showcase — NVIDIA released new physical-AI models alongside next-generation robots from a long roster of global partners, and the early-adopter list spans the humanoid pure-plays (Figure, Agility, Boston Dynamics), the logistics giants (Amazon Robotics), and industrial incumbents (Caterpillar, Hexagon).

But be skeptical and specific about what "design win" means here. A robot OEM running Isaac Sim for simulation or fine-tuning GR00T does not necessarily ship NVIDIA silicon in the production robot — some humanoid programs are evaluating Qualcomm's Dragonwing IQ10 as the power-efficient alternative for machines that do not need 2,000 teraflops, and Figure AI in particular has a publicly disclosed relationship with Qualcomm. The competitive reality is that NVIDIA owns the training-and-simulation layer almost uncontested, owns the high-end inference brain, but faces a real fight for the mid-tier and power-constrained inference sockets. The qualification timelines that matter run through 2026-2027 as the first humanoid programs move from prototype to pilot production — that is when "design win" converts to silicon revenue, or doesn't.

The competitive threat / Qualcomm

The named direct competitor on the robotics-compute layer is Qualcomm. At CES January 2026 Qualcomm launched the Dragonwing IQ10 series — an NPU rated at up to 700 TOPS, an 18-core Oryon CPU, support for 20-plus camera sensors and an integrated safety island — explicitly positioned as the power-efficient alternative to Jetson for robots that do not need Thor-class compute. Qualcomm has converted this into

commercial traction: a multi-year agreement with NEURA Robotics and a disclosed relationship with Figure AI, and management has gone as far as to say it expects robotics revenue to eventually surpass its entire smartphone-chip business (roughly $25 billion in FY25). That is an aggressive claim, but it signals Qualcomm is resourcing the fight seriously.

The bear-case-via-competitor is concrete: NVIDIA's data-center moat is CUDA plus scale, but in robotics the workload is inference at the edge, where power efficiency and cost matter more than raw FLOPS, and where Qualcomm's mobile-SoC heritage is a genuine advantage. NVIDIA's counter is the integrated stack — if GR00T becomes the default foundation model, the path of least resistance is to run it on NVIDIA silicon. There is no active IP litigation between the two. The realistic read is that robotics compute will be a two-horse race (with Ambarella a distant third in vision-specific sockets), NVIDIA taking the high-performance and training share, Qualcomm taking meaningful mid-tier inference share — which is fine for NVDA the stock, because robotics is not what the multiple rests on, but it caps how big the robotics call option can get.

The terminal risk

The terminal risk for NVDA is not robotics-specific — it is the custom-ASIC threat to the data-center business that actually sets the stock price. Every major hyperscaler is now designing in-house AI silicon (Google TPU, Amazon Trainium/Inferentia, Microsoft Maia, Meta MTIA), and the structural question is whether merchant GPUs lose share to custom accelerators for inference workloads over the 2027-2030 window. If that transition accelerates, NVDA's data-center growth decelerates hard, the multiple compresses, and the robotics call option — still ~1% of revenue — is nowhere near big enough to offset it. That is why the stock trades at only ~20x forward earnings despite the growth: the market is pricing the ASIC risk.

The robotics-specific terminal risk is milder but real: foundation models for robotics could commoditize, and the compute could fragment across Qualcomm, Ambarella, custom robot-OEM silicon and Chinese alternatives (Horizon Robotics and others) faster than NVIDIA can lock in the platform. NVIDIA's defense in both cases is the same — software lock-in and the integrated stack — and so far it has held. But a robotics investor in NVDA must accept that they are buying the data-center risk profile to get the robotics optionality.

Bull / Gap / Optionality (Photoncap framing)

Bull

1. NVIDIA owns the robotics software and simulation layer almost uncontested. Over 2 million developers use the NVIDIA robotics stack (company disclosure, August 2025), Isaac Sim and Isaac Lab are the de facto standard for sim-to-real training, and GR00T is the most-cited open humanoid foundation model. Even competitors' robots are often trained in NVIDIA's tools. If physical AI scales the way data-center AI did, NVIDIA has the equivalent of the CUDA moat already in place — and it monetizes through silicon attach.

2. The Jetson Thor silicon ladder is now complete and priced for volume. With AGX Thor ($3,499 dev kit, $2,999 T5000 module), the cheaper T4000 ($1,999), and the safety-certified IGX Thor, NVIDIA covers the full range from industrial AMRs to high-end humanoids. At $2,000-3,500 of compute content per robot, even a few million humanoid units a year — a plausible 2028-2030 number — is a multi-billion-dollar silicon line, and that is before the recurring software/Omniverse attach.

3. The early-adopter list is the best in the industry. Figure, Agility Robotics, Boston Dynamics, Amazon Robotics, Caterpillar, Hexagon, Medtronic and Meta are all building on the stack (NVIDIA, August 2025). Caterpillar's physical-AI partnership in particular signals that industrial heavyweights — not just startups — are standardizing on NVIDIA, which is where the durable volume eventually comes from.

4. The valuation is undemanding for the growth. At roughly 20.5x forward earnings against ~78% expected revenue growth in Q1 FY27 and a Strong Buy consensus with a ~$272 average target, the data-center business alone arguably justifies the price. That means the robotics franchise — Jetson, Isaac, GR00T, Omniverse — is effectively a free call option inside the stock, not something you are paying a premium for.

5. The cadence of model releases is relentless. GR00T N1 (March 2025) to N1.7 (March 2026) to a committed N2 by end-2026 on the new DreamZero architecture shows NVIDIA is iterating on robot foundation models faster than any competitor, and each release widens the gap in the software layer that ultimately pulls silicon.

Gap

1. Robotics is ~1% of revenue and the stock is set by data-center GPUs. The combined automotive-and-robotics segment is roughly $570-590 million a quarter against ~$68 billion total — about 1% on a reporting basis. Anyone buying NVDA "for robotics" is really buying data-center-GPU risk; the robotics thesis cannot move the stock for years, and a data-center disappointment swamps any robotics good news.

2. **The custom-ASIC threat to the core business is real and accelerating.** Google, Amazon, Microsoft and Meta are all ramping in-house AI silicon. If merchant-GPU share erodes for inference workloads over 2027-2030, NVDA's growth decelerates and the multiple compresses — and the ~20x forward P/E says the market already sees this risk. The robotics optionality is too small to offset it.

3. **Qualcomm is a credible competitor for the inference socket.** The Dragonwing IQ10 (CES January 2026, up to 700 TOPS) plus disclosed wins at NEURA and Figure show Qualcomm taking the power-efficient mid-tier seriously. NVIDIA's robotics silicon share is not the near-lock its data-center share is — and "design win on the software stack" does not guarantee NVIDIA silicon ships in the production robot.

4. **China export-control exposure remains a live swing factor.** The Q1 FY26 print carried a $4.5 billion H20 inventory charge and $2.5 billion of unshippable revenue. Policy on advanced-chip exports to China can move quarterly results by billions in either direction, and that volatility sits on top of the stock regardless of the robotics story.

Optionality

Event	Date / window	Direction
Q1 FY27 earnings	May 20, 2026 (after close)	Binary on FY27 data-center ramp; robotics commentary read-through
GR00T N2 release (DreamZero architecture)	End-2026 (committed)	Bull if it doubles task-success rates as claimed
First humanoid programs moving prototype → pilot production	2026-2027	Bull — converts design wins to Jetson silicon revenue
US-China advanced-chip export policy changes	Ongoing	Binary — multi-billion-dollar quarterly swings
Hyperscaler custom-ASIC ramp milestones	2027-2030	Bear — structural share-erosion risk to the core
Robotics broken out as a standalone reporting segment	Unconfirmed / if it happens	Bull — would signal robotics revenue has become material

The trade

NVDA is a Bucket B name for the robotics theme: a high-quality, reasonably-valued compounder where the robotics franchise is genuinely best-in-class but is a small call option on top of a data-center business that sets the price. Initiate in a **$221-244 entry zone** (current $232.77 minus roughly 1xATR to plus 5%), size at **2.0-3.0% of risk capital** — sizing this is really a data-center-AI allocation decision, not a robotics one — with a **stop near $198** (below the 50-day moving average and prior structural support). The defining near-term binary is **Q1 FY27 earnings on May 20, 2026**, where the print is about data center but the robotics commentary is the read-through this theme tracks. If you want a purer, higher-beta expression of the robotics-compute thesis without carrying the full data-center risk profile, Lattice Semiconductor (LSCC) is the cleaner sensor-fusion play and Ambarella (AMBA) the cleaner edge-vision play — both with robotics as a real growth vector rather than a 1% rounding error. NVDA's role in a robotics book is the anchor: the highest-quality name, the lowest thesis-specific conviction, owned as much for the platform optionality as for near-term robotics revenue. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

QCOM — QUALCOMM Incorporated · BUY (Tier-1) · Conv 7/10 · Bucket B

QUALCOMM Incorporated (QCOM)

The power-efficient robot brain — Dragonwing IQ10 SoCs that out-efficiency NVIDIA's Jetson, with disclosed humanoid wins at Figure and NEURA — bought at a single-digit-teens multiple funded by smartphone cash flow.

Investment Research · Photoncap-style deep dive · v1 of "Qualcomm" · May 14, 2026

What Qualcomm physically does

Qualcomm's robotics franchise is, at its core, the same thing that made it a $200-billion company applied to a new socket: a system-on-chip that integrates a CPU, an NPU, an image-signal processor, multi-camera interfaces and connectivity onto a single power-efficient die, plus the software stack to run AI workloads on it. In a smartphone that SoC is a Snapdragon; in a robot it is now branded Dragonwing. The mechanism is the same — Qualcomm's two-decade heritage is squeezing the most AI-and-vision compute out of the fewest watts, because phones run on batteries — and that heritage is exactly the constraint that matters in a robot, which also runs on a battery and cannot carry a data-center cooling solution.

The flagship robotics product is the Dragonwing IQ10 series, launched at CES January 2026. It is a full robot-brain SoC: an NPU rated at up to 700 TOPS, an 18-core Oryon CPU, support for more than 20 camera sensors, and an integrated functional-safety island. Below it sits the established Qualcomm Robotics RB3 platform — the RB3 Gen 2 development kit built on the QCS6490 chipset at roughly 12 TOPS, aimed at industrial AMRs, drones and lighter robots. The strategic framing Qualcomm uses is "physical AI from household robots up to full-size humanoids" — a single SoC family scaling across the whole robot size range.

Why is this a binding-constraint position rather than a me-too entry? Because the hard

problem in deploying humanoids at volume is not peak FLOPS — NVIDIA's Jetson Thor already wins on raw compute — it is running a capable vision-language-action model inside a power, thermal and cost budget that lets the robot be commercially viable. Qualcomm's pitch is that for the large majority of robots that do not need 2,000 teraflops, the IQ10 is the right answer: enough compute, far better efficiency, lower cost, and Qualcomm's mobile-grade manufacturing scale behind it. The honest framing, developed below, is that robotics sits inside Qualcomm's IoT segment and is not yet broken out as revenue — the robotics thesis on QCOM is that the company is the credible number-two in robot compute, bought at a value multiple because the market is focused on the declining handset business.

Product roadmap

The robotics line has two tiers. The established tier is the Qualcomm Robotics RB platform: RB3 Gen 2, built on the QCS6490 chipset at roughly 12 TOPS (a 4x AI-compute step over the prior RB3), with RB3 Gen 2 and RB3 Gen 2 Lite development kits shipping through partners such as Thundercomm and aimed at industrial, robotics and automation IoT use cases. This is the workhorse for AMRs, drones and lighter robots and has been in the market for the 2024-2025 robot generation.

The flagship tier launched at CES January 2026: the Dragonwing IQ10 series — up to 700 TOPS NPU, 18-core Oryon CPU, 20-plus camera sensors, integrated safety island — explicitly positioned for industrial AMRs and full-size humanoids. Alongside the silicon Qualcomm announced a full robotics stack: a comprehensive architecture integrating hardware, software and "compound AI." The branding consolidation matters too — Qualcomm folded its industrial and robotics edge-AI products under the Dragonwing brand through 2025-2026, and rolled out developer programs (MWC Talent Arena 2026, the AI Program for Innovators) to seed the ecosystem.

What Qualcomm does not make is the robot, the actuators, the sensors or a humanoid foundation model of its own — it partners on the model layer. It also still derives the majority of its revenue from smartphone SoCs and its licensing business, which is the context every robotics dollar has to be read against. The roadmap cadence is a roughly annual Snapdragon/Dragonwing generation refresh, with the IQ10 representing the deliberate push into humanoid-class compute.

The financial print

Qualcomm reported fiscal Q2 2026 on April 29, 2026: revenue of roughly $10.6 billion and non-GAAP EPS of $2.65, with EPS at the high end of guidance. The segment detail is where the robotics-adjacent story lives — automotive revenue surpassed a $5 billion annualized run-rate for the first time and was up 38% year-on-year, and QCT IoT (the segment that contains robotics) grew 9% year-on-year. The handset business, still the majority of QCT revenue, is the part the market worries about. On a full-year basis Qualcomm's FY25 revenue was in the high-$40-billion range; the smartphone-chip business alone was roughly $25 billion in FY25, which is the number Qualcomm's management has provocatively said it expects robotics revenue to eventually exceed.

The valuation is the headline: a forward P/E of approximately 19.1x, against a market cap of roughly $214.6 billion. That is a value multiple — the market is pricing the secular handset risk (Apple in-housing its modem, China OEM dynamics) and giving Qualcomm little credit for the automotive and IoT/robotics diversification. RSI 63.7 is mid-range and not extended, though the stock is +37.8% above its 50-day moving average, reflecting a strong recent run. Sell-side coverage is broad; the constructive case (a Motley Fool-cited bull thesis built on CEO Cristiano Amon's "the winner of edge AI wins the AI race" framing) leans on the diversification, while the bears stay focused on handset concentration.

The binary event is **Q3 FY26 earnings, expected late July 2026** (Qualcomm reports on a roughly quarterly cadence; July 30, 2026 is the working estimate, to be confirmed). The read-through this theme cares about is the IoT segment growth rate and any incremental disclosure on Dragonwing IQ10 customer traction.

Customer mix today

Qualcomm reports by segment, not by named customer, and robotics is not yet a disclosed line. The structure: QCT (the chip business) is dominated by handsets, with Android OEMs — Samsung and the Chinese majors — plus the Apple modem relationship that is in secular wind-down. Automotive is the fastest-growing reported segment, now above a $5 billion annualized run-rate and up 38% year-on-year (Q2 FY26). IoT, which contains industrial, edge and robotics, grew 9% year-on-year in Q2 FY26. The licensing business (QTL) is a separate, high-margin royalty stream.

The structural shift Qualcomm is selling: in 2023-2024 the QCOM story was "handset chip company with a litigation-prone licensing business and a declining Apple relationship." By 2026 the company is trying to reframe as "edge-AI compute company" — automotive above $5 billion, IoT growing, and robotics positioned as the next leg. On robotics customers specifically, Qualcomm has been more concrete than Ambarella: it

has disclosed a design win with Figure AI and a multi-year agreement with NEURA Robotics, both tied to the Dragonwing robotics push, plus a strengthened collaboration with Advantech on edge-AI and robotics applications. These are named, disclosed relationships — not yet sized as revenue, but more tangible than the typical "we have a pipeline" claim. The honest read: robotics is still inside the 9%-growing IoT bucket, but the named-customer disclosure suggests it is converting from narrative to backlog faster than the pure-plays.

What's actually happening at Figure and NEURA

The mechanism to watch is the humanoid-OEM SoC decision. Qualcomm's claim to be the credible number-two in robot compute rests on whether the best-funded humanoid programs actually design in the Dragonwing IQ10 rather than NVIDIA's Jetson Thor. The two disclosed proof points are Figure AI — where Qualcomm announced a design win — and NEURA Robotics, where Qualcomm signed a multi-year agreement. Management's framing on the Q2 FY26 call was that the IQ10 has generated "substantial customer interest" since the CES launch, with the NEURA agreement "reinforcing confidence in becoming a significant player in the broad robotics market."

Be specific and skeptical here. Figure AI is widely understood to also work with NVIDIA's software stack — a humanoid program evaluating both NVIDIA and Qualcomm silicon is the normal state of affairs at this stage, and "design win" can mean anything from a production socket to a development relationship. What Qualcomm has that NVIDIA's roster does not, in this specific comparison, is a credible efficiency-and-cost argument: for a humanoid that has to hit a commercially viable price point, the IQ10's 700-TOPS-at-mobile-power profile is a real alternative to Thor's 2,070-TFLOPS-at-130-watts. The qualification timelines run through 2026-2027 as Figure, NEURA and others move from prototype to pilot production — that is when the IQ10 either shows up in shipping robots or doesn't. The honest position: Qualcomm has the most concrete disclosed humanoid relationships of the merchant chip vendors, but they are not yet sized, and NVIDIA still owns the default.

The competitive threat / NVIDIA

The named competitor is NVIDIA, and the competitive dynamic is unusually clean. NVIDIA's Jetson Thor wins on raw compute (2,070 FP4 TFLOPS) and owns the integrated stack — Isaac for development, GR00T for the foundation model, Omniverse for

simulation — which creates the same kind of software gravity that CUDA created in the data center. Qualcomm's counter is the Dragonwing IQ10's efficiency and cost, plus a full robotics stack of its own, plus mobile-scale manufacturing. There is no active IP litigation between the two on the robotics front.

The bear-case-via-competitor is that NVIDIA's software lock-in pulls silicon: if GR00T becomes the default robot foundation model and it is optimized first and best for NVIDIA hardware, the path of least resistance for a robot OEM is Jetson, and Qualcomm is left fighting for the cost-sensitive tier. The bull-case-via-competitor is that the robot market is enormous and fragmented — household robots, industrial AMRs, drones, humanoids of every size — and a single integrated stack will not win all of it; Qualcomm's mobile heritage genuinely suits the high-volume, cost-constrained majority of that market. Qualcomm management's own framing — robotics revenue eventually exceeding the ~$25 billion smartphone-chip business — is an aggressive claim, but it signals the company is resourcing the fight as a strategic priority, not a side project. Ambarella is a distant third, competing only for the vision-specific socket.

The terminal risk

Qualcomm's terminal risk is not robotics-specific — it is whether the secular decline of the handset business outruns the diversification. Apple is in-housing its modem, which removes a multi-billion-dollar Qualcomm revenue stream over the next few years; the China handset market is structurally mature and the share dynamics among Chinese OEMs are volatile; and the QTL licensing business faces periodic renewal and litigation risk. If handsets decline faster than automotive, IoT and robotics grow, the whole company's revenue stalls and the ~19x multiple — which already prices skepticism — could still compress further. The robotics franchise, sitting inside the 9%-growing IoT segment, is not yet big enough to be the swing factor.

The robotics-specific terminal risk is the same socket-and-stack dynamic that hangs over every merchant chip vendor: if NVIDIA's integrated stack achieves CUDA-like lock-in in robotics, Qualcomm's addressable share is capped at the cost-sensitive tier, and the "robotics exceeds smartphones" ambition never materializes. Qualcomm's defense is genuine — efficiency, cost, scale, a full stack of its own, and a market large enough to support two winners — but a robotics investor in QCOM is buying the handset-decline risk profile to get the robotics optionality at a cheap price.

Bull / Gap / Optionality (Photoncap framing)

1. Qualcomm is the credible number-two in robot compute, bought at a value multiple. A forward P/E of roughly 19.1x is a single-digit-teens-style multiple for a company with a genuine claim to a top-two position in one of the largest structural compute markets being created. The robotics optionality is effectively free inside a stock priced for handset decline.

2. The disclosed humanoid wins are the most concrete among merchant chip vendors. A Figure AI design win and a multi-year NEURA Robotics agreement (both disclosed around the CES January 2026 Dragonwing launch) are named, specific relationships — more tangible than Ambarella's pipeline talk, and a real toehold against NVIDIA's default position.

3. The efficiency-and-cost argument is structurally sound. The Dragonwing IQ10's 700-TOPS-at-mobile-power profile is the right answer for the large majority of robots that do not need Jetson Thor's 2,070 TFLOPS. Qualcomm's two-decade heritage of maximizing compute-per-watt is exactly the constraint that binds in a battery-powered robot — this is not a stretch into an unfamiliar problem.

4. Automotive proves the diversification model works. Automotive crossed a $5 billion annualized run-rate and grew 38% year-on-year (Q2 FY26) — Qualcomm has already demonstrated it can take its SoC franchise into a non-handset vertical and scale it. Robotics is the same playbook applied again, and the automotive precedent de-risks the execution question.

5. Smartphone cash flow funds the build with no dilution. Unlike the small-cap robotics names that must raise capital to fund their roadmaps, Qualcomm self-funds the entire Dragonwing and robotics-stack effort from handset and licensing cash flow. The robotics call option costs the shareholder nothing in dilution.

Gap

1. Robotics is not yet a disclosed revenue line — it is inside 9%-growing IoT. For all the named wins, robotics is still a slice of the IoT segment that grew 9% year-on-year in Q2 FY26. The "robotics exceeds smartphones" framing is a multi-year ambition, not a current number, and the stock is set by handsets today.

2. The handset business is in secular decline. Apple in-housing its modem removes a major revenue stream; the China handset market is mature; QTL licensing faces renewal and litigation risk. If handsets decline faster than the diversification grows, the whole company stalls — and the robotics franchise is too small to offset it.

3. NVIDIA's integrated stack is the structural threat. If GR00T becomes the default

robot foundation model with software gravity like CUDA, Qualcomm's robot-compute share is capped at the cost-sensitive tier. "Design win" at Figure does not guarantee a production socket — the humanoid programs are evaluating both vendors.

4. The stock is up sharply and the value-multiple cushion is thinner than it looks. QCOM is +37.8% above its 50-day moving average — the easy re-rating may already be partly done. A handset-segment disappointment on the next print could compress the multiple regardless of robotics progress.

Optionality

Event	Date / window	Direction
Q3 FY26 earnings	~July 30, 2026 (to be confirmed)	Binary on IoT growth; Dragonwing traction read-through
Figure AI / NEURA programs moving to pilot production	2026-2027	Bull — converts disclosed wins to shipping silicon
Robotics broken out as a reported segment or sub-line	Unconfirmed / if it happens	Bull — signals robotics revenue is now material
Apple modem in-housing milestones	Ongoing through 2026-2027	Bear — quantifies the handset-decline drag
Additional named humanoid-OEM design wins	Ongoing	Bull

The trade

QCOM is a Bucket B name for the robotics theme: the credible number-two in robot compute, with the most concrete disclosed humanoid relationships of the merchant chip vendors, bought at a genuine value multiple because the market is fixated on the declining handset business. Initiate in a **$193-214 entry zone** (current $203.61 minus roughly 1xATR to plus 5%), size at **1.5-2.5% of risk capital** — the value multiple supports a slightly larger position than the extended small-caps, but the handset overhang argues for restraint — with a **stop near $178** (below the 50-day moving average and prior structural support). The defining near-term binary is **Q3 FY26 earnings, expected around July 30, 2026**, where the read-through is the IoT segment growth rate and

Dragonwing IQ10 customer commentary. If you want a purer robotics-compute expression without the handset-decline drag, NVIDIA (NVDA) is the higher-quality but more-fully-valued play; if you want the dedicated-vision angle, Ambarella (AMBA). QCOM's role in a robotics book is the value anchor — the cheapest credible way to own robot-compute optionality, owned because the downside is protected by the multiple and the upside is a genuine top-two position in a market being created. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

AMBA — Ambarella, Inc. · WATCH (Tier-2) · Conv 6/10 · Bucket C

Ambarella, Inc. (AMBA)

The pure-play edge-AI vision processor — one of only three Western public companies that can give a robot eyes — but the robotics line is still small behind automotive and security.

Investment Research · Photoncap-style deep dive · v1 of "Ambarella" · May 14, 2026

What Ambarella physically does

Ambarella makes the chip that does the seeing. An Ambarella SoC sits behind a camera — or behind several cameras — and turns raw image-sensor data into structured perception: it runs the image-signal-processing pipeline, then runs neural networks on that video to detect objects, estimate depth, segment scenes and fuse multiple camera feeds, all on-device at very low power. The company's defining technical asset is the CVflow architecture — a dataflow design that processes neural networks differently from a general-purpose GPU, trading flexibility for far better performance-per-watt on vision workloads specifically. The product families that carry this are the CV2, CV5, CV7 and CV3 lines (the CV3-AD family aimed at automotive autonomy), plus the higher-end N1 SoC built around 16 Arm Cortex-A78AE cores for multi-camera AI hubs, edge AI servers and industrial robotics.

The reason this matters for robotics is that perception is the front of the pipeline. A humanoid or an autonomous mobile robot needs to interpret its environment before it can plan or act, and doing that vision processing on a power-constrained machine is exactly the problem CVflow was built for. Ambarella's pitch is a clean division of labor: NVIDIA trains the models in the cloud and increasingly wants the inference brain too, but Ambarella competes for the dedicated vision-inference socket — the part of the robot that handles cameras specifically — where performance-per-watt beats raw FLOPS.

The honest framing, developed below, is that robotics is not yet Ambarella's business. The company is a vision-chip company whose revenue today is automotive ADAS/electronic-mirror/driver-monitoring and enterprise video security, with AIoT,

industrial and robotics as the fastest-growing but still-minority slice. The robotics thesis on AMBA is that the CVflow architecture is genuinely well-suited to robot vision, and that as humanoids and AMRs scale, the dedicated-vision-SoC socket becomes a real market — and Ambarella is one of only three credible Western public vendors (alongside NVIDIA and Qualcomm) positioned for it.

Product roadmap

The CVflow generations form the spine. CV2, CV5 and CV7 are the established edge-AI vision families — management has explicitly named CV2, CV5 and CV7 as the key revenue drivers for fiscal 2027. The CV7 was showcased at CES January 2026 as the latest general-purpose edge-AI perception SoC. On the automotive side, the CV3-AD family targets ADAS and autonomous-driving central compute. At the top of the range, the N1 SoC — positioned for multi-camera AI hubs, edge AI servers running multi-modal LLMs, and industrial robotics — is the part most directly aimed at the robotics opportunity, paired with an 8K vision SoC launched in January 2026 for high-resolution multi-sensor perception.

The strategic repositioning is visible in how Ambarella markets itself: at Embedded World 2026 (February 2026) the company's theme was "The Ambarella Edge: From Agentic to Physical AI," and it maintains a dedicated AIoT/Industrial/Robotics product line. The cadence is roughly an annual refresh of the CVflow families on advancing process nodes, with the N1 and 8K parts representing the deliberate push up-market into robot-and-server territory.

What Ambarella does not make is the rest of the robot brain — it does not make the general-purpose application processor that runs planning and control (that is NVIDIA Jetson or Qualcomm territory), it does not make actuators or sensors, and it does not build foundation models. It is a vision-perception specialist, and its robotics thesis depends on the dedicated-vision socket remaining a distinct, defensible part of the bill of materials rather than being absorbed into an integrated SoC.

The financial print

Ambarella closed fiscal 2026 (year ended late January 2026) with revenue of $390.7 million, up 37.2% year-on-year — a record. The most recent quarterly print, Q4 FY26 reported February 25, 2026, delivered revenue of $100.9 million, up 20.1% year-on-year and ahead of the roughly $100.3 million consensus, with non-GAAP gross margin of

59.8% (the midpoint of the 59-60.5% guided range) and non-GAAP EPS of $0.13, a 30% beat. Full-year non-GAAP net profit was $26.9 million, or $0.62 per diluted share. The edge-AI segment drove the year, up roughly 50% year-on-year and contributing about 80% of full-year revenue. Cash and marketable securities at quarter-end were $312.6 million against a roughly $3.59 billion market cap.

Guidance for Q1 FY27 is revenue of $97-103 million. The forward P/E of roughly 75.5x is the number to sit with: this is a company whose multiple already prices in years of edge-AI growth, and at RSI 78.3 and +37% above its 50-day moving average it is technically extended — not in the AMBQ/MRAM league, but firmly in chase territory. Sell-side coverage is thinner than the mega-caps; the Zacks consensus framework and the usual edge-AI-focused analysts (Needham, Craig-Hallum, Rosenblatt have historically covered it) are the visible voices, generally constructive on the edge-AI mix shift but wary of the valuation.

The binary event is **Q1 FY27 earnings, expected May 26, 2026** (Ambarella reports roughly a month after quarter-end; date to be confirmed by company schedule). The read-through this theme cares about is any incremental disclosure on the AIoT/industrial/robotics mix and N1 traction.

Customer mix today

Ambarella does not publish customer names as revenue percentages — the disclosure is by end-market. The structural picture: management has guided that by the end of fiscal 2026, automotive and enterprise security together would account for over 85% of revenue. Within that, automotive is OEMs and Tier-1 suppliers building ADAS, electronic mirrors, driver- and cabin-monitoring and autonomous-driving systems; enterprise security is the professional video-surveillance integrators. That leaves roughly 15% or less for AIoT, industrial and robotics combined — the fastest-growing slice but still a minority.

The shift that is the actual story: edge AI went from a modest contributor a few years ago to roughly 80% of fiscal 2026 revenue, growing about 50% year-on-year. That is the mix shift Ambarella bulls are paying 75x forward earnings for — the company transitioning from a legacy video-compression chip vendor (its origins were action cameras and broadcast) into an edge-AI-perception company. Robotics specifically is inside the AIoT/industrial bucket, and while management talks about it more every quarter, it is honest to say robotics is an emerging application for Ambarella, not yet a reported revenue driver. The 2024-to-2026 change worth highlighting is the edge-AI share climb; the robotics share is still small enough that the company does not size it.

What’s actually happening in robotics adoption

The mechanism to watch is the N1 socket. Ambarella’s robotics traction, such as it is, runs through the N1 and 8K vision SoCs being designed into multi-camera robot perception hubs and into the AI processing for autonomous fleet and industrial-robotics applications. The company markets the N1 explicitly for “implementing industrial robotics” and “edge AI servers running multi-modal LLMs,” which is the on-device-vision pitch — let the application processor handle planning while the Ambarella part handles the camera array.

But be specific and skeptical about where this stands. Ambarella has not disclosed named humanoid-robot design wins the way Qualcomm has disclosed Figure and NEURA, or the way NVIDIA has disclosed its early-adopter roster. The robotics exposure is real at the architecture level — CVflow is genuinely good at robot vision — but at the revenue level it is still pipeline, not print. The third-party read (an “AMBA gets bought, and humanoids are the reason” thesis circulating among edge-AI investors) frames Ambarella as a potential acquisition target precisely because it owns a defensible vision-IP block that a larger robotics-SoC player might want — which tells you the market sees the value as strategic and optionality-driven rather than as a current robotics revenue stream. The qualification timelines that would convert this run through 2026-2027 as robot programs lock their bills of materials.

The competitive threat / NVIDIA and Qualcomm

The competitive set is unusually clear: there are exactly three relevant Western public companies that ship the kind of edge-AI chips a robot needs — NVIDIA, Qualcomm and Ambarella — and NVIDIA owns the current humanoid-silicon conversation through its tightly integrated Jetson/Isaac/GR00T stack. That is the core competitive problem. NVIDIA and Qualcomm both sell a full robot-brain SoC; Ambarella sells a vision-perception block. If robot OEMs decide they want one integrated SoC that does perception, planning and control together — which is the direction NVIDIA’s Jetson Thor and Qualcomm’s Dragonwing IQ10 both point — then the dedicated-vision socket Ambarella depends on gets absorbed, and Ambarella is left selling into the shrinking set of designs that still partition vision separately.

In automotive, the named competitor is Mobileye, plus the broader ADAS-SoC field; in security and computer-vision chips, a long tail of emerging challengers. There is no

active IP litigation that bears on the robotics thesis. The bear-case-via-competitor is concrete and quantifiable: NVIDIA's robotics-plus-automotive segment alone is roughly $570-590 million a quarter — larger than Ambarella's entire annual revenue of $390.7 million — and NVIDIA is iterating its stack faster. Ambarella's defense is performance-per-watt on vision specifically and the flexibility of being a merchant vision-IP supplier that does not compete with the robot OEM's own ambitions — a real but narrow moat.

The terminal risk

The terminal risk is socket absorption. The entire Ambarella robotics thesis rests on the dedicated edge-vision SoC remaining a distinct line item in the robot's bill of materials. If the industry standardizes — as NVIDIA and Qualcomm are pushing it to — on a single integrated robotics SoC that handles vision, planning and control on one die, then a vision-only chip vendor is structurally squeezed: it either moves up into full-SoC territory (where it would be a sub-scale competitor against NVIDIA and Qualcomm) or it accepts being a component supplier into a shrinking share of designs. The transition timing is the 2027-2030 window as humanoid bills of materials standardize. The named beneficiaries of that transition are precisely NVIDIA and Qualcomm.

The milder version of the same risk shows up in Ambarella's core markets: in automotive, ADAS central-compute is consolidating onto fewer, larger SoCs; in security, the camera-chip market is mature. So even setting robotics aside, Ambarella's structural challenge is that "the dedicated vision chip" is a category under pressure from integration. The company's credible-roadmap answer is the N1 — pushing up into multi-camera-hub and edge-server territory where a powerful dedicated vision processor still makes sense — but whether that holds against integrated SoCs is the open question that caps the multiple.

Bull / Gap / Optionality (Photoncap framing)

1. Ambarella is one of only three credible Western public edge-AI vision vendors. With NVIDIA and Qualcomm the only peers, Ambarella has scarcity value — there is no fourth Western public pure-play, which is why the market treats it as a strategic asset and potential acquisition target. In a robotics theme, owning the only listed pure-play vision-perception chip vendor has portfolio-construction value beyond the fundamentals.

2. The edge-AI mix shift is real and fast. Edge AI grew roughly 50% year-on-year in fiscal 2026 and is now about 80% of revenue (company disclosure, Q4 FY26). The

company has genuinely transformed from a legacy video-compression vendor into an edge-AI-perception company, and that transition is what justifies a growth multiple even before robotics contributes.

3. The CVflow architecture is genuinely differentiated for robot vision. The dataflow design delivers materially better performance-per-watt on neural-network vision workloads than a general-purpose GPU — the exact constraint that matters in a power-limited robot. If robot OEMs do partition vision separately, Ambarella is the natural merchant supplier, and the N1 SoC is purpose-built for that role.

4. The financials are clean and the balance sheet is strong. Record FY26 revenue of $390.7 million up 37.2%, non-GAAP gross margin near 60%, non-GAAP profitability, and $312.6 million of cash against a $3.59 billion cap. This is not a cash-burning story-stock — it is a profitable, growing company where robotics is upside, not a survival bet.

5. The strategic-acquisition optionality is a real floor. The "AMBA gets bought" thesis among edge-AI investors is not idle — Ambarella owns a defensible vision-IP block that a larger robotics-SoC player (or an automotive Tier-1) might want, which puts a takeout floor under the stock that the standalone fundamentals alone would not.

Gap

1. Robotics is still pipeline, not revenue. Automotive and security are over 85% of revenue (FY26 guidance); AIoT/industrial/robotics combined is the remaining ~15%, and robotics is only a slice of that. There are no disclosed named humanoid design wins. Anyone buying AMBA "for robotics" is buying an edge-AI-vision company with a robotics call option, not a robotics company.

2. The valuation is extended on both fundamentals and tape. A forward P/E of roughly 75.5x prices in years of growth, and at RSI 78.3 and +37% above the 50-day moving average the stock is in chase territory. Any disappointment on the edge-AI growth rate — let alone robotics — leaves significant multiple-compression room.

3. Socket absorption by integrated SoCs is the structural threat. NVIDIA's Jetson Thor and Qualcomm's Dragonwing IQ10 both push toward a single integrated robot-brain SoC. If that wins, the dedicated-vision socket Ambarella depends on shrinks — and NVIDIA's robotics-plus-automotive segment alone already out-revenues all of Ambarella.

4. Customer concentration and end-market cyclicality. Automotive is the largest end-market and is cyclical; enterprise security is mature. The company's growth depends heavily on the edge-AI upgrade cycle within those markets continuing — a macro auto downturn or a security-capex pause would hit the print regardless of the robotics narrative.

Optionality

Event	Date / window	Direction
Q1 FY27 earnings	~May 26, 2026 (to be confirmed)	Binary on edge-AI growth rate; robotics-mix read-through
First named humanoid/AMR design win disclosure	2026-2027	Bull — converts robotics pipeline to credible revenue path
N1 / 8K SoC design-in milestones in robotics hubs	2026-2027	Bull if disclosed
Strategic acquisition of Ambarella	Unscheduled / ongoing speculation	Bull — takeout premium
Integrated robotics-SoC standardization (NVIDIA/Qualcomm)	2027-2030	Bear — socket-absorption risk

---

The trade

AMBA is a Bucket C name for the robotics theme: a genuinely differentiated, profitable edge-AI-vision pure-play, but trading at ~75x forward earnings and +37% above its 50-day line, with robotics still a call option rather than a revenue stream. Initiate cautiously in a **$78-86 entry zone** (current $81.96 minus roughly 1xATR to plus 5%) — the tape is extended enough that scaling in on weakness is the disciplined approach — size at **1.0-2.0% of risk capital** given the valuation and the still-speculative robotics exposure, with a **stop near $67** (below the 50-day moving average and prior consolidation). The defining near-term binary is **Q1 FY27 earnings, expected around May 26, 2026**, where the read-through is the AIoT/industrial/robotics mix and N1 traction. If you want a cleaner expression of the same edge-AI thesis with deeper liquidity and a lower multiple, Qualcomm (QCOM) is the better-capitalized robotics-SoC play; if you want higher beta on the FPGA-sensor-fusion angle, Lattice (LSCC). AMBA's role in a robotics book is the pure-play vision satellite position — owned for the scarcity value and the strategic-takeout optionality as much as for near-term numbers, and sized accordingly. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

AMBQ — Ambiq Micro, Inc. · SKIP / WAIT (Tier-3) · Conv 4/10 · Bucket D

Ambiq Micro, Inc. (AMBQ)

Ultra-low-power edge-AI silicon with a real physical-AI optionality leg — but a wearables-and-IoT core, persistent losses, and the single most overbought tape in the entire research universe.

Investment Research · Photoncap-style deep dive · v1 of "Ambiq Micro" · May 14, 2026

Two-layer framing — read this first

This name requires the two-layer discipline before anything else, because the chart and the story are pulling in opposite directions. **Layer one — the tape:** AMBQ prints an RSI of 89.6 and sits +115.4% above its 50-day moving average. Those are not "strong momentum" readings — they are among the most extreme overbought readings in the entire coverage universe, the kind of extension that historically mean-reverts hard regardless of how good the story is. **Layer two — the thesis:** the underlying physical-AI/edge-AI story has a genuine kernel of merit, and robotics is a real (if emerging) application. The job of this deep dive is to hold both layers at once: take the thesis seriously, and refuse to let the thesis override the entry-timing discipline. The conclusion, developed in full below, is that AMBQ is a Bucket D name — the near-term move has played out, the fundamentals do not yet justify the price, and the only defensible posture is a watch-and-wait stance with a deep-pullback entry zone, not a chase. Anyone tempted to buy this at $72 because "robotics" should re-read this paragraph.

What Ambiq physically does

Ambiq makes microcontroller-class and SoC-class silicon whose entire reason for existing is power: it runs AI workloads at a fraction of the energy a conventional chip would consume. The technical asset is the SPOT architecture — Subthreshold Power Optimized Technology — a design approach that operates transistors at or below the

threshold voltage where they normally switch, squeezing out dramatic energy savings at the cost of design complexity. The product family that carries this is the Apollo line of ultra-low-power SoCs. The result is a chip that can run neural-network inference — keyword spotting, gesture recognition, health-signal analysis, sensor anomaly detection — on a device that has to live for days or weeks on a small battery.

Ambiq's historical home is wearables: smartwatches, fitness bands, hearables — devices where battery life is the product and an Apollo SoC running on-device AI is the enabling component. The company has reported that over 80% of shipped units run AI algorithms, which is the genuine differentiation: this is an edge-AI company, not a generic MCU vendor.

Where does robotics fit, and how honestly? Robotics is an emerging application, not the core market. The robot connection is real at the architecture level — a humanoid or an AMR is full of always-on, power-constrained sensor subsystems (touch sensors, joint encoders, environmental monitors, presence detection) that need to run inference continuously without burdening the main battery or the main brain, and that is precisely the SPOT/Apollo use case. But the dollars today are wearables, with the fastest-growing adjacent markets being medical, industrial and smart home/buildings. Robotics sits inside that industrial slice. The physical-AI thesis on AMBQ is that as sensor-rich machines proliferate, ultra-low-power edge inference becomes a pervasive requirement and Ambiq's subthreshold IP is a durable advantage — a plausible long-term claim, but one whose robotics-specific revenue is not yet visible in any disclosure.

Product roadmap

The roadmap is the Apollo SoC generations built on the evolving SPOT architecture. The Apollo family has progressed through multiple generations of ultra-low-power SoCs, each adding AI-inference capability — the company's framing is "higher-value Apollo platforms" driving volume and average-selling-price up as customers adopt the more capable, more AI-centric parts. Ambiq's Q1 2026 commentary emphasized that the mix shift toward these higher-value platforms is what is driving revenue growth and the gradual gross-margin improvement.

Beyond the silicon, Ambiq's roadmap is really a market-expansion roadmap: moving the SPOT/Apollo franchise out of wearables and into medical, industrial and smart-home/building applications, with robotics as a sub-segment of industrial. Management has guided that medical, industrial and smart home/buildings are each expected to deliver revenue growth exceeding 100% in 2026, and that non-wearables are now roughly a quarter of the pipeline. As a recent IPO (Ambiq listed in July 2025), the company's

public-market roadmap is also a profitability roadmap — management has explicitly said it needs roughly $47 million of quarterly revenue to reach breakeven, with breakeven targeted for late 2027 or early 2028.

What Ambiq does not make is a robot brain, an application processor, a foundation model, or any robotics-specific product line — it makes general ultra-low-power edge-AI silicon that robotics can use, the same way wearables and medical devices use it. There is no "robotics SoC" in the Ambiq catalog. That is an important honesty point: the robotics exposure here is application pull, not a dedicated product.

The financial print

Ambiq reported Q1 2026 on May 12, 2026: net sales of $25.06 million, up 59.3% year-on-year and 20.8% sequentially, driven by edge-AI demand and the mix shift to higher-value Apollo platforms. GAAP gross margin was 43.5%, non-GAAP gross margin 46.2%. The company posted a GAAP net loss of $10.2 million and a non-GAAP net loss of $5.0 million, or $0.25 per share — this is an unprofitable company, which is why the forward P/E field is not meaningful. Non-wearable revenue grew 100% year-on-year and non-wearables are now about a quarter of the pipeline.

Guidance for Q2 2026 is net sales of $31.0-32.0 million, non-GAAP gross margin of 45.0-46.0%, and a non-GAAP net loss per share of $0.23-0.29 — so losses continue, and on a widening absolute basis as the company invests. CFO commentary put the breakeven bar at roughly $47 million of quarterly revenue, with breakeven targeted for late 2027 or early 2028. Market cap is roughly $1.54 billion against a company doing ~$100 million of annualized revenue and losing money — that is a price-to-sales multiple in the mid-teens for an unprofitable small-cap, which is the valuation context for the overbought tape.

The binary event is **Q2 2026 earnings, expected mid-August 2026** (Ambiq reports on a roughly quarterly cadence; August 11, 2026 is the working estimate, to be confirmed). For a name this extended, the next print is a genuine risk event in both directions — a beat-and-raise could extend the move further, but any wobble against a +115%-above-50MA setup invites a violent unwind.

Customer mix today

Ambiq does not disclose named customers as revenue percentages — the disclosure is by end-market. The structural picture: wearables have historically been the majority of revenue, the legacy core where smartwatch and fitness-band makers use Apollo SoCs for on-device AI and battery life. The shift the company is selling is diversification away from that concentration: non-wearables (medical, industrial, smart home/buildings combined) are now roughly 25% of the pipeline and grew 100% year-on-year in Q1 2026, with each of those three sub-markets guided to >100% growth in 2026.

The 2024-to-2026 change worth highlighting: Ambiq is transitioning from a wearables-dependent component vendor — a structurally concentrated, consumer-cyclical position — toward a broader edge-AI silicon supplier. Robotics specifically is inside the industrial sub-segment of that non-wearables 25%, which means robotics is, at most, a single-digit-percentage slice of revenue today and Ambiq has not sized it. The honest framing: the diversification is real and is happening in the numbers (non-wearables doubling), but robotics as a distinct revenue contributor is not yet a disclosed line — it is a sub-component of a sub-segment, and the robotics narrative around AMBQ is running well ahead of the robotics revenue.

What’s actually happening in the non-wearables push

Because Ambiq has no named humanoid or AMR design wins to point to, the mechanism to watch is the non-wearables ramp broadly — medical, industrial, smart-home — with robotics as one industrial thread inside it. The Q1 2026 evidence: non-wearable revenue grew 100% year-on-year, roughly a quarter of the pipeline is now non-wearables, and order activity included expedited-delivery requests that management read as healthy underlying demand. Over 80% of shipped units run AI algorithms, which tells you the customer base is genuinely buying Ambiq for edge inference, not just for a cheap MCU.

Be specific and skeptical about robotics within this. Ambiq has not disclosed a single named robotics customer, a robotics design-win count, or a sized robotics pipeline. There is no Ambiq equivalent of Qualcomm’s “Figure design win” or NVIDIA’s early-adopter roster. The robotics exposure is inferential: robots contain many always-on low-power sensor subsystems, that is the Apollo use case, therefore as robots scale Ambiq benefits. That logic is sound but it is a thesis about the future, not a description of the present. The qualification-to-revenue path for any robotics-specific contribution runs out past 2026 — and in the meantime the name’s revenue and the name’s stock price are being driven by wearables-plus-medical-plus-general-industrial, with robotics as a story attached. For a name trading at a mid-teens price-to-sales multiple with RSI 89.6, that gap between narrative and disclosed fact is the central risk.

The competitive threat / the MCU-plus-NPU squeeze

Ambiq's competitive threat is the broad microcontroller and SoC industry adding AI capability. The named competitors are the large MCU and edge-SoC vendors — STMicroelectronics, NXP, Renesas, Infineon, Texas Instruments, and the broader Arm-Cortex-M ecosystem — all of which are adding neural-processing units and AI-acceleration blocks to their mainstream MCU lines. Ambiq's differentiation is the subthreshold SPOT architecture delivering genuinely better energy-per-inference, but the question is whether that advantage is wide enough and durable enough to defend a sub-scale, unprofitable specialist against giants with vastly more design resources, broader catalogs, deeper customer relationships and the ability to bundle.

The bear-case-via-competitor is concrete: Ambiq does roughly $100 million of annualized revenue and loses money; STMicroelectronics, NXP and TI each do tens of billions and are profitable. If a "good enough" ultra-low-power AI MCU from one of those vendors closes most of the energy gap, Ambiq's addressable advantage narrows to a niche, and a niche does not support a $1.5 billion market cap. There is no IP litigation defining the competitive timeline — this is a slow-grinding share-and-relevance question, not a courtroom event. Ambiq's defense is that subthreshold operation is genuinely hard to replicate and that for the most extreme power-constrained applications it remains the best answer — plausible, but "best at the extreme niche" is a smaller business than the current valuation implies.

The terminal risk

The terminal risk is category absorption. Ultra-low-power edge AI as a distinct silicon category exists only as long as general-purpose MCU and SoC vendors cannot deliver "good enough" energy efficiency with their mainstream parts. As STMicroelectronics, NXP, Renesas, TI and the Arm ecosystem keep adding NPUs and refining low-power modes, the gap that justifies a dedicated subthreshold specialist narrows. The transition window is continuous and already underway — every MCU generation from the majors chips away at it. If the gap closes faster than Ambiq can scale to profitability and entrench in defensible sockets, a sub-scale unprofitable specialist gets squeezed between "good enough from the giants" above and commodity pricing below. The named beneficiaries of that transition are precisely the large MCU incumbents.

Ambiq's credible defense is that subthreshold operation is a genuinely differentiated,

patent-protected approach and that the most demanding always-on edge-AI applications will always reward the lowest-energy solution. There is a real business in being the specialist at the extreme. But “real business” and “$1.5 billion market cap at a mid-teens price-to-sales multiple while losing money” are different statements, and the terminal risk is less that Ambiq disappears and more that it never grows into — or even toward — its current valuation. That is the constraint that, combined with the overbought tape, drives the Bucket D classification.

Bull / Gap / Optionality (Photoncap framing)

1. The SPOT subthreshold architecture is a genuine, hard-to-replicate differentiation. Operating transistors at or below threshold voltage delivers real energy-per-inference advantages that the majors have not matched, and over 80% of Ambiq’s shipped units run AI algorithms — this is a real edge-AI company with real IP, not a story stock with no product.

2. The diversification away from wearables is working in the numbers. Non-wearable revenue grew 100% year-on-year in Q1 2026, non-wearables are now ~25% of the pipeline, and medical/industrial/smart-home are each guided to >100% growth in 2026 — the concentration risk is genuinely declining, not just being talked about.

3. The physical-AI long-term logic is sound. Sensor-rich machines — including robots — are full of always-on, power-constrained inference subsystems, which is exactly the Apollo use case. If physical AI scales the way bulls expect, pervasive ultra-low-power edge inference is a large structural market and Ambiq’s IP is positioned for it.

4. Revenue growth is accelerating off a small base. 59.3% year-on-year growth in Q1 2026 with Q2 guided to a further sequential step — the top-line trajectory is steep, and if it sustains, the path to the ~$47 million quarterly breakeven bar by late 2027/early 2028 is at least visible.

5. As a recent IPO it is under-owned by institutions and lightly covered. Ambiq listed only in July 2025; the float dynamics and thin sell-side coverage that contribute to the violent tape can also cut the other way — genuine fundamental progress can re-rate a name like this fast once institutional coverage builds. (This is a double-edged point, included honestly as a bull driver because it is part of why the name moves.)

Gap

1. The tape is the single most extended in the coverage universe. RSI 89.6 and +115.4% above the 50-day moving average are extreme-mean-reversion readings.

Buying here is buying at the most dangerous point of the move — the entry-timing risk alone is disqualifying for a full position regardless of the thesis.

2. The company loses money and will keep losing money. A non-GAAP net loss of $5.0 million in Q1 2026, Q2 guided to a further loss, and breakeven not expected until late 2027 or early 2028 at the earliest. The valuation — roughly $1.5 billion cap on ~$100 million annualized revenue — prices in flawless execution toward a profitability target two years out.

3. Robotics is narrative, not disclosed revenue. No named robotics customers, no robotics design-win count, no sized robotics pipeline. Robotics is a sub-thread of the industrial sub-segment of the 25%-of-pipeline non-wearables bucket — the robotics story is running far ahead of the robotics dollars.

4. The category-absorption terminal risk is active now. STMicroelectronics, NXP, Renesas, TI and the Arm ecosystem are all adding NPUs to mainstream MCUs every generation. The gap that justifies a dedicated subthreshold specialist is narrowing continuously, and Ambiq is sub-scale and unprofitable while the squeeze plays out.

5. The core market is consumer-cyclical wearables. Despite the diversification, wearables are still the revenue majority — a consumer-electronics demand pause would hit the print hard, and the diversification is not yet far enough along to insulate the company.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~August 11, 2026 (to be confirmed)	Binary — and a violent one given the extension
First named robotics design-win disclosure	Unscheduled / past 2026	Bull — would convert narrative to fact
Non-wearables crossing 50% of revenue	2027 (if growth sustains)	Bull — de-risks the concentration
Reaching the ~$47M quarterly breakeven run-rate	Late 2027 / early 2028 (management target)	Bull — removes the going-concern discount
Major MCU vendor closing the ultra-low-power AI gap	Ongoing	Bear — terminal-risk acceleration

The trade

AMBQ is a Bucket D name: the near-term move has, by any technical measure, already played out, and the fundamentals — an unprofitable company two years from breakeven, with robotics still a narrative rather than a revenue line — do not justify chasing the price here. The disciplined posture is watch-and-wait. The **entry zone is set deliberately deep at $54-62** — roughly 15-25% below the current $72.25 — because the two-layer framework demands it: an RSI-89.6, +115%-above-50MA name with no near-term fundamental catalyst to defend the price is exactly the profile where the entry zone should be a real pullback target, not "current ± 5%." Size at **0.5% of risk capital at most**, and only on a genuine reversion into that zone — this is a starter/probe position in a speculative, unprofitable small-cap, not a core holding. Set a **stop near $47** (below the deep-entry zone, accepting that a name this volatile needs room). The catalyst that matters is **Q2 2026 earnings around August 11, 2026** — but note that for this name the next print is as much a risk event as an opportunity. If you want exposure to the ultra-low-power edge-AI-in-robotics thesis without the going-concern risk and the parabolic tape, there is no clean listed pure-play substitute — the better-capitalized way to own edge-AI-for-robotics is QCOM or LSCC, and AMBQ should be understood as the high-risk, deep-pullback-only lottery ticket of this batch. **Conviction: 4 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

MRAM — Everspin Technologies, Inc. · SKIP / WAIT (Tier-3) · Conv 4/10 · Bucket D

Everspin Technologies, Inc. (MRAM)

The dominant pure-play MRAM vendor — the instant-state-save memory inside factory robots — with a genuine industrial-automation core and a new defense deal, but a tiny, lumpy P&L on the most extended tape in this batch.

Investment Research · Photoncap-style deep dive · v1 of "Everspin Technologies" · May 14, 2026

Two-layer framing — read this first

This name demands the two-layer discipline up front, because the tape and the story diverge sharply. **Layer one — the tape:** MRAM sits +173.7% above its 50-day moving average with an RSI of 78.0. The +173.7% figure is among the most extreme extensions in the entire coverage universe — a near-tripling relative to the 50-day mean. Extensions like that are momentum, but they are also the setups that mean-revert most violently when the catalyst flow pauses. **Layer two — the thesis:** the underlying business has a genuine kernel — Everspin is the real dominant pure-play in a real (if small) memory category, industrial automation is a confirmed core end-market, and the recently announced $40 million defense subcontract is a tangible, dated piece of news, not vapor. The job here is to take that kernel seriously while refusing to let it excuse chasing a parabolic chart. The conclusion, built out below, is Bucket D: the move has played out, the fundamentals are too small and too lumpy to support the current price, and the only defensible posture is a deep-pullback watch stance. Anyone tempted to buy MRAM near $39 because "robots need state-save memory" should re-read this paragraph first.

What Everspin physically does

Everspin makes MRAM — magnetoresistive random-access memory — a non-volatile memory technology that stores each bit in the magnetic orientation of a tiny structure called a magnetic tunnel junction rather than in an electrical charge. The practical consequences of that physical difference are the whole investment case: MRAM is non-volatile (it keeps its data with the power off, like flash), but it is also fast and effectively endless in write-endurance (like SRAM or DRAM), and it is radiation-tolerant. Everspin ships two families — Toggle MRAM, the established, extremely robust generation, and STT-MRAM (spin-transfer torque), the higher-density modern generation — plus it has announced the UNISYST MRAM family aimed at the high-density standalone NOR-flash-replacement market.

Why does this matter for robotics specifically? The killer application in a factory robot or an industrial automation system is instant state-save on power loss. When a programmable logic controller, a robot arm controller or an automated machine loses power — a fault, an e-stop, a brownout — it has milliseconds to write its current operating state somewhere non-volatile so it can resume cleanly rather than restart from zero, potentially mid-motion, in a way that could damage product or hurt a person. Conventional flash is too slow to write in that window; battery-backed SRAM works but adds a battery that ages and fails. MRAM writes fast enough to capture the full machine state in the available window and holds it with no battery. That is why MRAM has a genuine, defensible home in industrial automation — and as factories deploy more robots, the attach grows. The honest caveat, developed below, is that this is a small, niche, lumpy memory market, and Everspin's whole company does roughly $60 million of annual revenue — the robotics-and-automation thesis is real at the use-case level but the scale is tiny.

Product roadmap

Everspin's roadmap runs across three product directions. The established core is Toggle MRAM — the extremely robust, well-qualified generation that has been the company's bread and butter in industrial, automotive and aerospace/defense applications for years, valued for reliability and radiation tolerance over density. The modern density leg is STT-MRAM, spin-transfer-torque MRAM, the generation that competes on capacity and is qualified into automotive and industrial reliability standards on 22nm and 28nm flows. The newest direction is the UNISYST MRAM family, which Everspin has said is expected to expand its addressable market by roughly $3 billion by targeting high-density standalone NOR-flash replacement — a deliberate push beyond the company's niche-state-save heartland into a much larger memory category.

Alongside the product families, the strategic-roadmap event of 2026 is the defense sub-

contract: in its Q1 2026 disclosures Everspin announced a new 2.5-year, $40 million agreement under which it acts as a subcontractor on an existing prime contract, providing Toggle MRAM process-technology capabilities and engineering services for U.S. defense-industrial-base customers. That is a roadmap item as much as a revenue item — it deepens Everspin's defense entrenchment, where MRAM's radiation tolerance is genuinely differentiated.

What Everspin does not make is mainstream memory — it does not make DRAM, NAND flash or the high-volume commodity memory that fills phones and PCs and data centers. It is a specialty-memory company. Its roadmap risk, addressed in the terminal-risk section, is that the foundries (TSMC, GlobalFoundries, Samsung) offer embedded MRAM as IP within their own process design kits, which is a different go-to-market than Everspin's standalone-chip model.

The financial print

Everspin reported Q1 2026 on April 29, 2026: total revenue of $14.9 million, up 14% year-on-year and at the high end of guidance. The composition matters — MRAM product sales were $14.1 million, up 28% year-on-year, while the licensing, royalty and patent line fell to $0.8 million from $2.1 million in the prior-year quarter, which is the lumpiness in action. GAAP gross margin was 52.7%, up from 51.4%, and non-GAAP net income was $2.6 million, or $0.11 per diluted share. Management attributed the strength to industrial automation, transportation and data-center applications. Guidance for Q2 2026 is revenue of $15.5-16.5 million, with GAAP EPS guided to a net loss of $0.07-0.12 and non-GAAP EPS from breakeven to a $0.03 profit.

Put the scale in perspective: this is a company with roughly $60 million of annualized revenue and a market cap of about $918 million — a price-to-sales multiple around 15x and a forward P/E around 84.2x, for a business growing its product line in the mid-20s-percent range but with a total top line measured in tens of millions and a profit line that flickers between small profit and small loss quarter to quarter. The $40 million defense subcontract spread over 2.5 years is about $16 million of cumulative revenue — meaningful against a $60 million base, which is exactly why the stock reacted, but it does not transform the company into something the current valuation comfortably fits.

The binary event is **Q2 2026 earnings, expected late July 2026** (Everspin reports on a roughly quarterly cadence; July 29, 2026 is the working estimate, to be confirmed). For a name +173.7% above its 50-day line, the next print is a high-stakes risk event — anything short of continued acceleration against that setup invites a sharp unwind.

Customer mix today

Everspin discloses by end-market application, not by named customer percentage. The structural picture from the Q1 2026 disclosures: the MRAM product business is driven by industrial automation, transportation and data center, with industrial automation the core and the most relevant to the robotics theme. Defense is a fourth and now-expanding leg via the $40 million subcontract. Everspin has not published a "Customer A is X%" breakout, and given the small revenue base, customer concentration is a real but undisclosed risk — a memory company doing $15 million a quarter very likely has meaningful exposure to a handful of large industrial and defense accounts.

The 2024-to-2026 change worth highlighting: the mix is tilting toward the higher-value, higher-margin applications. MRAM product sales grew 28% year-on-year in Q1 2026 while the lower-quality licensing/royalty line shrank — so the revenue is becoming more product-driven and arguably higher-quality, even as the absolute scale stays tiny. Within "industrial automation," factory robotics and automated machinery are a genuine and growing slice, but Everspin does not size robotics as a distinct line. The honest framing: industrial automation as a category is confirmed core and contains real robotics demand; robotics specifically is a sub-component the company does not break out; and the customer base is concentrated enough — though undisclosed — that the lumpy quarter-to-quarter revenue pattern is partly a concentration symptom.

What's actually happening in industrial automation and defense

The mechanism to watch is twofold. On the industrial-automation side, the driver is the secular build-out of factory automation — every additional robot arm, every additional PLC, every additional automated machine is a potential MRAM socket for instant state-save, and Everspin's Q1 2026 commentary explicitly named industrial automation as a strength. This is the cleanest robotics read-through: more factory robots, more state-save memory, and Everspin is the dominant pure-play supplier of it. The catch is that this is a slow-compounding, GDP-plus kind of growth, not a hypergrowth ramp — it shows up as the mid-20s-percent MRAM product growth, not as a step-change.

On the defense side, the $40 million, 2.5-year subcontract announced with Q1 2026 results is the concrete event. Everspin acts as a subcontractor to a U.S. prime, supplying Toggle MRAM process technology and engineering services for defense-industrial-base customers. This is where MRAM's radiation tolerance is genuinely irreplaceable — mission-critical defense, satellite and aerospace systems need memory that survives ra-

diation and holds state, and that is structurally Everspin's strongest competitive position. Be specific and skeptical, though: the prime contractor is unnamed, the $16-million-or-so cumulative revenue is spread thinly over 2.5 years, and a single subcontract — however welcome — does not by itself justify a near-tripling of the stock relative to its 50-day average. The defense deal is real and good; the market's reaction to it is the part that requires discipline. The qualification-and-ramp timelines for both legs are gradual — this is not a business with a single dated catalyst that re-rates it; it is a slow grind that the tape has gotten far ahead of.

The competitive threat / the foundries and Avalanche

Everspin's competitive picture has two distinct fronts. The first is other MRAM suppliers. Avalanche Technology is the named direct pure-play competitor, and in March 2026 Avalanche reported progress scaling STT-MRAM magnetic-tunnel-junction cells toward 64Gb-128Gb space-grade products for defense, satellite and aerospace — directly contesting Everspin's strongest niche. The much larger structural competitors are the foundries: TSMC, GlobalFoundries and Samsung all offer embedded MRAM (eMRAM) as IP within their process design kits, so a chip designer who wants MRAM-like non-volatile memory can increasingly get it embedded on a foundry process rather than buying a standalone Everspin chip. The top of the broader STT-MRAM market by the market-research tallies is Samsung, TSMC, SK Hynix, Micron and Intel — giants for whom MRAM is a feature, not a company.

The bear-case-via-competitor is structural: Everspin is the dominant pure-play, but "pure-play" in a category that the foundries can offer as embedded IP is a precarious position. Everspin's defense is genuine — it holds deep MRAM-specific process IP, has the longest qualification track record (which matters enormously in industrial, automotive and defense), and the standalone-chip model still wins where a designer needs MRAM but is not building on a foundry process that offers eMRAM. There is no IP litigation defining the competitive timeline. But the competitive reality caps the multiple: Everspin can be the best pure-play MRAM company in the world and still be a sub-$100-million-revenue niche supplier, because the addressable market for standalone MRAM is itself bounded.

The terminal risk

The terminal risk is that MRAM stays a permanent niche — and that the standalone-pure-play slice of that niche gets squeezed from two sides. From below, conventional memory keeps getting cheaper and "good enough": for many applications, slower flash plus a supercapacitor, or battery-backed SRAM, remains the cost-competitive answer to state-save, and MRAM's premium is only justified where the timing and battery-free constraints are hard. From above, embedded MRAM on foundry processes (TSMC, GlobalFoundries, Samsung eMRAM offerings) absorbs the designs being built on those processes, removing them from Everspin's standalone-chip addressable market. The transition is not a single dated event — it is the slow, continuous evolution of where MRAM gets sourced. The named beneficiaries are the foundries and the conventional-memory incumbents.

Everspin's credible defense is that there is a durable, defensible core — industrial, automotive, aerospace and defense applications where qualification track record, radiation tolerance and the battery-free state-save guarantee are non-negotiable, and where Everspin's standalone parts and process IP genuinely win. The UNISYST family is the attempt to break out of the niche into the larger NOR-flash-replacement market, which is the bull's growth path. But the honest terminal-risk read is that MRAM has been "the memory of the future" for two decades and remains a small category — the risk is not that Everspin fails, it is that the company stays small and lumpy indefinitely while trading at a valuation that requires it to break out. That, combined with the +173.7% tape, is what drives the Bucket D classification.

Bull / Gap / Optionality (Photoncap framing)

1. Everspin is the genuine dominant pure-play in a real memory category. It is the recognized MRAM frontrunner with the broadest Toggle-and-STT portfolio and the longest qualification track record — in industrial, automotive and defense, where qualification history is decisive, that incumbency is a real moat within the niche.

2. Industrial automation is a confirmed core end-market with a real robotics read-through. Management explicitly named industrial automation as a Q1 2026 strength, and the instant-state-save use case in factory robots and automated machinery is genuine and grows as factory automation scales — this is the cleanest robotics-thesis link in the name.

3. The $40 million defense subcontract is concrete, dated and entrenching. A 2.5-year subcontract supplying Toggle MRAM process technology to the U.S. defense industrial base deepens Everspin's position exactly where MRAM's radiation tolerance

is irreplaceable — and at roughly $16 million cumulative against a ~$60 million revenue base, it is materially additive.

4. The revenue mix is becoming higher-quality. MRAM product sales grew 28% year-on-year in Q1 2026 while the lower-quality licensing/royalty line shrank — the business is becoming more product-driven, with GAAP gross margin improving to 52.7%.

5. The UNISYST family is a credible attempt to break out of the niche. Everspin's stated ~$3 billion addressable-market expansion via standalone-NOR-flash-replacement is the growth optionality — if UNISYST gains traction, the "permanent small niche" terminal risk is partly answered.

Gap

1. The tape is +173.7% above the 50-day moving average — among the most extreme in the universe. A near-tripling relative to the 50-day mean is a setup that mean-reverts violently when catalyst flow pauses. Buying here is buying at the most dangerous point of a parabolic move — the entry-timing risk alone disqualifies a full position.

2. The business is tiny and lumpy. Roughly $60 million of annualized revenue, a profit line that flickers between small profit and small loss quarter to quarter, and a licensing line that swung from $2.1 million to $0.8 million year-on-year. A ~$918 million market cap and ~84x forward P/E on that base prices a breakout that has not happened.

3. The foundry-eMRAM and conventional-memory squeeze is structural. TSMC, GlobalFoundries and Samsung offer embedded MRAM as IP; cheaper flash-plus-supercapacitor and battery-backed SRAM remain "good enough" for many state-save needs. Everspin's standalone-pure-play addressable market is bounded from both sides.

4. Customer concentration is a real, undisclosed risk. A memory company doing ~$15 million a quarter almost certainly leans on a handful of large industrial and defense accounts; Everspin does not disclose the breakout, and the lumpy revenue pattern is partly a concentration symptom.

5. MRAM has been "the memory of the future" for two decades and stayed niche. The terminal risk is not failure — it is permanent smallness. The valuation requires a breakout the category's twenty-year history has not delivered.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~July 29, 2026 (to be confirmed)	Binary — and high-stakes given the +173.7% extension
UNISYST MRAM family gaining design traction	2026-2027	Bull — addresses the permanent-niche terminal risk
Additional defense subcontracts or prime-contract wins	Ongoing	Bull
Foundry eMRAM offerings broadening	Ongoing	Bear — standalone-market erosion
Avalanche Technology scaling space-grade STT-MRAM	2026-2027	Bear — direct competition in Everspin's strongest niche

---

The trade

MRAM is a Bucket D name: the near-term move has, on every technical measure, played out, and the fundamentals — a ~$60 million-revenue, lumpy, niche memory company trading at ~84x forward earnings — do not justify chasing a stock sitting +173.7% above its 50-day line. The disciplined posture is watch-and-wait. The **entry zone is set deliberately deep at $29-34** — roughly 13-26% below the current $39.15 — because the two-layer framework requires it: a name this extended with no near-term catalyst to defend the price is precisely where the entry zone must be a real reversion target, not "current ± 5%." Size at **0.5% of risk capital at most**, and only on a genuine pullback into that zone — a probe position in a tiny, volatile specialty-memory name, never a core holding. Set a **stop near $26** (below the deep-entry zone, with room for the volatility this name carries). The catalyst is **Q2 2026 earnings around July 29, 2026** — which, for a name this extended, is as much a risk event as an opportunity. There is no cleaner listed pure-play substitute for the MRAM-in-robotics thesis — Everspin is the pure-play — so the honest pivot is that the broader, cleaner way to own factory-automation-driven robotics demand is through the automation and compute names in the theme rather than through a sub-$100-million-revenue memory niche on a parabolic chart. MRAM's role in a robotics book is a deep-pullback-only watch item, not an initiate-now position. **Conviction: 4 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

Section 7 — Layer 5 — OEMs: where competition fragments

The OEM layer is where everything above gets integrated into a finished, sellable robot. Structurally it is the worst layer in the stack — OEMs buy from concentrated suppliers at oligopoly pricing, integrate at moderate margin, and compete with each other on price. Most are loss-making or pre-profit. The strategic-vehicle category is the cheapest exposure (robot programs inside larger profitable businesses); the humanoid pure-plays are the most concentrated bet (and the least profitable). The industrial-robotics incumbents are the "own the existing market and option the new one" call. Specialty mobile robotics is the wildcard category.

Four sub-tiers within the OEM layer:

(a) Humanoid pure-plays. UBTECH (9880.HK) is the only listed pure-play actually shipping full-size humanoids into customer factories (FY2025 revenue ~RMB 2.0B, +53% YoY, humanoid mix 41%, still loss-making at RMB 790M operating loss). Unitree's planned Shanghai STAR IPO (A-share — not IBKR-tradable, excluded from this basket) is the world's top humanoid seller by unit volume and uniquely profitable. Agibot's planned Hong Kong IPO Q3 2026 (HKEX — would be IBKR-tradable) is the most-funded Chinese competitor; add on listing. Western pure-plays — Serve Robotics (SERV), Richtech (RR), Nauticus (KITT) — are smaller and more speculative, with Nauticus in distressed territory post-reverse-split.

(b) Strategic-vehicle OEMs. Robot programmes sit inside larger profitable businesses, giving optionality at a cheaper multiple. Tesla (TSLA — Optimus, V3 prototype Q1 2026, target 50,000-100,000 unit ramp in 2026 with the Fremont retooling). Hyundai Motor (005380.KS — owns 80% of Boston Dynamics, the cleanest listed Atlas/Spot proxy, at 14× earnings). Rainbow Robotics (277810.KQ — Samsung-controlled 35% post-March 2025 acquisition, Korean industrial conglomerate optionality). This is the cheapest way to own humanoid exposure in our universe.

(c) Industrial robotics incumbents. Fanuc (6954.T), Yaskawa Electric (6506.T) and ABB are the established profitable industrial robot leaders, where humanoid is option

value on top of a $20-80B industrial robot annual demand base. ABB has the most concrete near-term catalyst: a planned Q2/H2 2026 robotics-division separation (Q1 2026 commentary now references mid-2026 close).

(d) Specialty mobile and warehouse robotics. Symbotic (SYM — warehouse robotics systems for Walmart, the only GAAP-profitable name in the developer bucket, with a $22B+ backlog and Walmart concentration of ~84%). Teradyne (TER — owns Universal Robots and MiR, but robotics is ~7% of group revenue; the rest is semiconductor test). Serve Robotics (autonomous sidewalk delivery). Richtech and Nauticus, as noted, are the speculative tail.

005380 — Hyundai Motor (005380.KS) · BUY (Tier-1) · Conv 7/10 · Bucket B

Hyundai Motor (005380.KS)

The only profitable, cash-generative way to own a top-tier humanoid program — Atlas, vertically integrated, on an automaker's balance sheet.

Investment Research · Photoncap-style deep dive · v1 of "Hyundai Motor" · May 14, 2026

What Hyundai physically does — from the robotics lens

This is not a generic automaker write-up. Hyundai Motor is covered here because it owns roughly 80% of Boston Dynamics — the company that builds Atlas, widely regarded as the most technically advanced humanoid robot platform in the world — and because Hyundai is doing something no other humanoid player is: building a vertically integrated humanoid robotics supply chain inside a single industrial conglomerate, funded by the cash flow of a profitable car business. Atlas is the product that matters. The new generation of Atlas, unveiled by Boston Dynamics, is a fully electric humanoid (the prior hydraulic Atlas was a research platform; the electric one is built for production) designed, in Boston Dynamics' own words, to be "the most production-friendly robot we've ever designed" — engineered with reduced unique part counts and components compatible with automotive supply chains.

The reason the automotive ownership is the whole thesis: a humanoid robot's hardest problems are mass-manufacturing reliable actuators at cost, and having a real customer to deploy into. Hyundai solves both structurally. Hyundai Mobis — the group's tier-1 components arm — agreed at CES 2026 to supply actuators for Atlas; actuators convert control signals into motion and represent more than 60% of a humanoid's material cost, so whoever industrialises actuator production controls humanoid economics. Hyundai Mobis brings automotive-grade mass-manufacturing and reliability engineering to that problem. And Hyundai's own factories are the first deployment ground: Atlas's first commercial deployments in 2026 go to Hyundai's Robotics Metaplant Application Center (RMAC) for automotive manufacturing, with the entire 2026 Atlas production run committed internally to RMAC and to Google DeepMind for AI research.

The core auto business is the engine that funds all of this. Hyundai Motor is one of the world's largest automakers — over 4.1 million vehicles sold in 2025 — generating the operating cash flow that pays for a $26 billion US investment program that includes a dedicated robotics factory. The robotics exposure is an embedded, self-funded option inside a profitable, cash-generative industrial company. That structure — not Atlas's joint articulation — is what makes Hyundai unique in this theme.

Product roadmap — the robotics roadmap

The product roadmap, robotics-side, runs through Atlas and the manufacturing build-out. The new electric Atlas was unveiled by Boston Dynamics as the production-intent generation. The confirmed deployment milestones: Atlas's first commercial deployments in 2026 are at Hyundai's Robotics Metaplant Application Center (RMAC) and at Google DeepMind; Hyundai has committed its entire 2026 Atlas production run to those two destinations — internal use and one strategic AI-research partner, no open-market sales yet. Atlas was named "Best Robot" in CNET's Best of CES 2026 awards, an external validation point.

The manufacturing roadmap is the structurally important part. Hyundai Motor Group announced (originally August 2025, reaffirmed at CES 2026) plans for a US robotics factory targeting 30,000 Atlas units per year by 2028, as part of a $26 billion US investment program. Hyundai Mobis's CES 2026 agreement to supply Atlas actuators is the first customer in Mobis's new robotics-components business and the spine of the vertically integrated supply chain. Atlas deployment at Hyundai Motor Group Metaplant America in Savannah, Georgia is targeted by 2028, beginning with proven-safety processes like parts sequencing.

What is confirmed: the new Atlas exists, 2026 deployments to RMAC and Google DeepMind are committed, the 30,000-unit/year factory is announced with a 2028 target, and Mobis is the actuator supplier. What is roadmap/intent rather than confirmed: that the 30,000-unit capacity will actually fill with economically justified demand, that external Atlas customers beyond Google DeepMind materialise, and that the 2028 timelines hold. Treat the build-out as a credibly funded plan, not a delivered fact — the binding question is demand, not capability.

The financial print

Hyundai Motor reported FY2025 results (announced late January 2026) of record revenue of KRW 186.3 trillion (~$132B USD), up 6.3% year-on-year, but operating profit of KRW 11.47 trillion (~$8.1B USD), down 19.5%, for a 6.2% operating margin. Net profit was KRW 10.36 trillion, down 21.7%. The profit decline was driven by US tariffs and higher incentives in an intensely competitive market — Q4 2025 alone saw operating profit down 39.9% YoY. This is the crucial point for a robotics investor: Hyundai is being valued at a forward P/E of just 14.3 *because* the auto business is under tariff and competition pressure — and that depressed multiple means you are getting the Boston Dynamics / Atlas optionality close to free, embedded inside a stock the market is pricing as a struggling automaker.

For FY2026, Hyundai guided to consolidated revenue growth of 1–2% and an operating margin of 6.3–7.3% — i.e., a modest stabilisation, with the tariff drag still present but the company targeting 4.16 million-plus units. The robotics spend sits inside the $26 billion US investment program; it is funded from auto cash flow and a strong balance sheet, not from a dilutive raise or debt stress.

The binary event ahead is the Q2 2026 earnings, expected around July 23, 2026 — the print that shows whether the auto-business margin is stabilising per guidance and whether management gives any updated colour on the robotics build-out timeline. The 1-year stock return into May 2026 has been strong — the stock is +34.8% above its 50-day MA with RSI at a stretched 78.6 — so the auto-recovery-plus-robotics-optionality trade is no longer undiscovered, and the entry has to respect that the easy part of the move has happened.

Customer mix today

Hyundai’s customer mix has to be read on two layers. The auto business — the cash engine — sells over 4.1 million vehicles a year to a globally diversified retail customer base across North America, Europe, Korea, India and emerging markets; no customer-concentration risk, the classic diffuse automaker mix. That diversified, profitable, cash-generative base is what funds the robotics bet, and it is the reason Hyundai is a fundamentally different risk profile from any pre-profit humanoid pure-play.

The robotics customer mix is, deliberately, almost entirely captive today. The first and dominant Atlas customer is Hyundai itself — RMAC, then Metaplant America by 2028. The only disclosed external Atlas customer for the 2026 production run is Google DeepMind, taking units for AI research rather than production deployment. Boston Dynamics also has its established Spot quadruped business with an external industrial-inspection

customer base, but Atlas — the humanoid that drives the thesis — has a 2026 customer mix of "Hyundai internal plus one AI-research partner." On the components side, Hyundai Mobis's first robotics-components customer is Boston Dynamics (Atlas actuators) — a captive intra-group relationship at the start. The structural shift to watch is the transition from captive-only to genuine external Atlas demand: the day Boston Dynamics signs an Atlas customer outside the Hyundai group and outside research use is the day the robotics business becomes a real franchise rather than an internal automation program. That has not happened yet.

What's actually happening at RMAC and inside the Hyundai supply chain

The mechanism to watch is the vertical integration loop, and it is genuinely differentiated. Hyundai's plan is: Boston Dynamics designs Atlas for manufacturability; Hyundai Mobis industrialises the actuators (60%+ of material cost) using automotive-grade mass-production and reliability discipline; Hyundai's own factories (RMAC first, Metaplant America by 2028) are the deployment ground and the proving environment; and the 30,000-unit/year US factory scales production. Each piece reinforces the others — Mobis gets a guaranteed first customer to learn actuator mass-production on, Boston Dynamics gets a reliable low-cost supplier and a captive deployment site to generate operating data, and Hyundai gets automation it controls end to end. Boston Dynamics explicitly cited Mobis's "reliability-based evaluation systems and global-standard mass production capabilities" as the reason for the partnership.

This is the closest thing in the entire humanoid theme to a real industrialisation flywheel, and it is why Hyundai screens as the highest-quality robotics exposure in this Asian batch. But the honest mechanism check: as of May 2026, the loop is being *built*, not *running at scale*. RMAC is an "Application Center" — a proving environment — not a 30,000-unit production line. The 30,000-unit factory targets 2028. Atlas's 2026 deployment is committed but small. And the question the flywheel cannot yet answer is the demand question: even with the best actuator supply chain in the world, the 30,000-unit capacity only makes sense if Atlas units actually pay for themselves on a factory floor versus fixed automation — and that economic proof does not yet exist. The vertical integration de-risks the *cost and supply* side of humanoids better than anyone; it does not de-risk the *demand* side. RMAC's job through 2026–2027 is to generate exactly that demand-side proof.

The competitive threat / Tesla Optimus and the well-funded field

Hyundai/Boston Dynamics' competitive set on the humanoid frontier is the most credible group in the theme: Tesla Optimus (Tesla's own captive-deployment-first humanoid, with Tesla's manufacturing scale and balance sheet behind it — the closest analog to Hyundai's strategy, an industrial company building a humanoid for its own factories first), the well-funded US humanoid startups (Figure and others, venture-funded, racing on AI and deployment partnerships), and the Chinese field led by UBTECH and Unitree (lower cost structure, China industrial-policy tailwind, UBTECH already #1 in full-size humanoid shipments). Rainbow Robotics, Hyundai's Korean peer, is also in the frame but is sub-scale and behind Boston Dynamics on bipedal capability.

The competitive argument for Hyundai is that it has the best *combination* of assets: Boston Dynamics' technology lead on Atlas, an automaker's balance sheet and cash flow, Mobis's actuator manufacturing, and captive deployment sites. Tesla has a similar structural model but is the more direct competitor for exactly that reason. The Chinese players are ahead on shipped units and cost but behind on the kind of vertically integrated, automotive-grade supply chain Hyundai is constructing. The bear competitive case: Atlas's technical lead may not translate into a commercial lead if cost-competitive Chinese humanoids "good enough" the market first, and a 30,000-unit factory is a large bet that could be stranded if a competitor's economics prove out faster. There is no material IP litigation involving Hyundai/Boston Dynamics in the humanoid space as of May 2026. The competitive risk is not legal — it is that being the most-advanced and most-integrated player still does not guarantee being the player whose unit economics work first.

The terminal risk

The terminal risk for Hyundai-as-robotics-play is the humanoid category's universal terminal risk, refracted through a large capital commitment: humanoids stay uneconomic versus fixed automation, and the 30,000-unit/year factory becomes stranded capex. Hyundai is committing real money — the robotics factory sits inside a $26 billion US investment program — on the bet that general-purpose humanoids reach task-economics parity. If that parity slips into the 2030s, Hyundai has built a humanoid factory that cannot fill, and it has done so while the auto business that funds everything is being ground down by US tariffs and price competition. The compounding risk is the double exposure: the terminal humanoid risk and the cyclical-plus-structural auto risk hitting at the same time.

The transition timing is the same embodied-AI maturation curve that governs the whole

theme — parity in 2027–2029 validates the build; parity in the 2030s strands it. The mitigant, and it is a real one, is that Hyundai's downside is bounded in a way no pure-play's is: if humanoids disappoint, Hyundai is still a profitable automaker trading at 14x earnings, and the robotics capex is a manageable fraction of a $26B program — the stock does not go to zero, it just loses the optionality premium. The named alternative beneficiaries of a "humanoids stay uneconomic" world are the industrial-arm incumbents (Fanuc, Yaskawa) and, as ever, the components supply chain — though note that in Hyundai's case, the components supplier *is* inside the group (Mobis), which means even a partial humanoid success has a path to value via Mobis. Hyundai has the most credible roadmap for the transition of any name in this batch; the multiple you can pay is constrained not by roadmap doubt but by the auto-business cyclical drag and the sheer size of the capacity bet.

Bull / Gap / Optionality

1. The only profitable, cash-generative way to own a top-tier humanoid program. Hyundai prints KRW 11.47 trillion of operating profit and trades at a 14.3x forward P/E. Every other credible humanoid bet — UBTECH, Rainbow, Doosan, the US startups — is loss-making or pre-revenue. Hyundai lets you own Atlas optionality on a profitable, dividend-paying balance sheet, with the downside bounded by a real auto business.

2. The vertical integration loop is the theme's best industrialisation flywheel. Boston Dynamics designs Atlas for manufacturability, Hyundai Mobis industrialises the actuators (60%+ of humanoid material cost) with automotive-grade discipline, Hyundai's factories deploy and generate operating data, and a 30,000-unit/year US factory scales it. No competitor — except arguably Tesla — has all four pieces inside one group.

3. The robotics optionality is close to free at today's multiple. The market is pricing Hyundai as a tariff-pressured automaker at 14x. Atlas — named "Best Robot" at CES 2026, with committed 2026 deployments and a funded factory build — is embedded inside that multiple at little to no premium. You are not paying a humanoid valuation; you are paying an automaker valuation and getting the option.

4. Atlas is the technology leader, externally validated. The new electric Atlas is widely regarded as the most advanced humanoid platform, designed explicitly for production, and won CNET's Best of CES 2026 "Best Robot" award. Hyundai owns ~80% of the company that builds it. The technology lead is real, not narrative.

5. Captive deployment solves the humanoid "first customer" problem. Atlas's 2026 production is fully committed to RMAC and Google DeepMind; Hyundai's own

factories are a guaranteed deployment ground and data-generation environment. A humanoid program's biggest early risk — no one to sell to — is structurally solved by Hyundai having millions of square feet of its own factory floor.

Gap

1. **The auto business that funds everything is under tariff pressure.** FY2025 operating profit fell 19.5%, Q4 fell 39.9%, on US tariffs and incentive competition. FY2026 guidance is for only 1–2% revenue growth. The cash engine is being ground down, and if it deteriorates further the robotics spend competes for constrained capital.

2. **The 30,000-unit factory is a large bet on unproven demand.** The 2028 capacity target only pays off if Atlas units are economically justified on a factory floor versus fixed automation — and that proof does not exist yet. RMAC is an "Application Center," not a running production line. This is stranded-capex risk on a multi-billion-dollar commitment.

3. **The robotics customer mix is captive — zero proven external Atlas demand.** Atlas's 2026 mix is Hyundai-internal plus Google DeepMind (research). No external production customer has been signed. Until Boston Dynamics sells Atlas outside the Hyundai group for real deployment, this is an internal automation program, not a franchise.

4. **The stock is technically stretched into the entry.** RSI 78.6, +34.8% above the 50-day MA — the auto-recovery-plus-robotics-optionality trade is well discovered and the stock has run hard. The risk of a 10–15% mean-reversion pullback before the next leg is elevated, and chasing at these levels has real near-term drawdown risk.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~Jul 23, 2026	Binary on auto-margin stabilisation + robotics colour
First external Atlas production customer signed	2026–2027	Bull — converts internal program to franchise
US robotics factory construction/capacity milestones	2026–2028	Bull on progress; bear if delayed

Event	Date / window	Direction
RMAC deployment data — Atlas economics on a real line	2026–2027	Binary — the demand-side proof
US tariff policy developments	Ongoing	Binary — drives the auto cash engine both ways
Hyundai Mobis actuator-business scaling disclosures	2026–2027	Bull — validates the vertical integration loop

---

The trade

Hyundai Motor is the highest-conviction name in this Asian humanoid batch, precisely because it is the one where you are not betting the company on humanoids — you are buying a profitable, cash-generative automaker at 14x earnings and getting a top-tier, vertically integrated humanoid program (Atlas, via ~80%-owned Boston Dynamics) as an embedded option close to free. The discipline issue is entry: the stock is technically stretched (RSI 78.6, +34.8% above its 50-day MA), so initiate at KRW 676,400–747,600 (current KRW 712,000 ± 5%) and be willing to leg in — take a starter position now and add on any 5–10% mean-reversion pullback, which the price-discipline playbook expects in this kind of run. Size at 1.5% of risk capital — the largest sizing in this batch, justified by the bounded downside: if humanoids disappoint, you still own a 14x-earnings automaker, not a zero. Stop at KRW 600,000, roughly 16% below current, beneath the structural breakout level — the auto business plus the embedded optionality means a deeper stop is acceptable because the thesis is not fragile. The catalyst is the Q2 2026 earnings around July 23, 2026, the read on auto-margin stabilisation and any robotics build-out colour; the longer-horizon catalysts are the first external Atlas customer and RMAC deployment economics. If you want a purer, higher-beta expression of the same vertically integrated humanoid thesis, Hyundai Mobis is the components-layer play — the actuator supplier monetising 60%+ of Atlas's material cost — and on the house view that the supply chain monetises a robotics theme before the OEMs do, Mobis is the cleaner shot, with Hyundai Motor the lower-risk, profitable anchor. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

6506 — Yaskawa Electric (6506.T) · BUY (Tier-1) · Conv 7/10 · Bucket B

Yaskawa Electric (6506.T)

Industrial robots plus the dominant servo-motor franchise — Yaskawa makes the motion "muscle" inside every robot, humanoids included.

Investment Research · Photoncap-style deep dive · v1 of "Yaskawa Electric" · May 14, 2026

What Yaskawa physically does

Yaskawa Electric is one of the two Japanese industrial-automation giants alongside Fanuc, but its centre of gravity is different and, for the robotics theme, arguably better positioned. Yaskawa runs two core businesses. The first is industrial robots — the Motoman line of six-axis articulated arms used for welding, handling, painting and assembly, where Yaskawa is one of the global "big four" robot makers. The second, and the one that makes Yaskawa structurally interesting for this theme, is Motion Control: AC servo motors and the drives that control them, plus inverters. Yaskawa is one of the dominant servo-motor makers in the world.

Here is why the servo position is the key to the thesis. A servo motor is the precision actuator — the "muscle" — that makes any robot move accurately: it converts an electrical command into controlled rotation, with the position feedback that lets a robot arm hit a sub-millimetre target. Every industrial robot, every cobot, and critically every humanoid robot is, at the hardware level, a collection of servo or actuator joints under coordinated control. Yaskawa does not just make complete robots; it makes the motion-control technology that goes *inside* robots — including, potentially, the humanoids that are trying to disrupt the industrial-robot business. That dual position — sell the finished robot, and sell the muscle that powers everyone else's robot — is what makes Yaskawa a more theme-resilient name than a pure robot-arm maker.

The third leg is a smaller systems-engineering business. Yaskawa does not make its own semiconductors or AI foundation models, but it has moved aggressively to embed AI into its products: the strategic direction in 2025–2026 is "physical AI," pairing Yaskawa's

motion hardware with AI from NVIDIA and a social-scale deployment partnership with SoftBank. Yaskawa is positioning itself as the motion layer of the embodied-AI era, not just a legacy robot vendor.

Product roadmap

The Motoman industrial-robot line is the established core — six-axis arms across the payload range for automotive and general-industry automation. The roadmap event that re-rates Yaskawa for the robotics theme is the MOTOMAN NEXT platform: a next-generation robot platform built for "autonomous adaptivity," with an open development environment so robots can execute complex tasks autonomously in unstructured environments. MOTOMAN NEXT comprises five industrial robots (the NEX4, NEX7, NEX10, NEX20 and NEX35) and two collaborative robots (the NHC12 and NHC30), and — critically — they are powered by NVIDIA Isaac accelerated libraries and AI models running on NVIDIA Jetson Orin autonomous control units. This is Yaskawa explicitly building the bridge from "programmed industrial robot" to "AI-driven adaptive robot."

The strategic-partnership roadmap is the other half. In November 2025, Yaskawa and SoftBank Corp. announced a strategic collaboration to implement "Physical AI" at social scale — combining Yaskawa's MOTOMAN NEXT autonomous robots with SoftBank's AI-RAN communication infrastructure for real-time, low-latency robot decision-making, targeting robots for offices, hospitals and schools. Yaskawa is also integrating NVIDIA Omniverse and Isaac simulation frameworks into its virtual-commissioning solutions and putting NVIDIA Jetson modules into its controllers for edge AI inference.

On the Motion Control side, the roadmap is steady iteration of servo motors, drives and inverters — the franchise that feeds semiconductors, machine tools and EV production. What Yaskawa does *not* have is a finished, named bipedal humanoid product with a launch date — the SoftBank "physical AI" partnership is about autonomous robots for human environments, and Yaskawa's natural humanoid play is as a motion-component and platform supplier rather than a humanoid OEM. Treat MOTOMAN NEXT and the SoftBank/NVIDIA partnerships as confirmed; treat any "Yaskawa humanoid OEM" expectation as not on the public roadmap.

The financial print

Yaskawa Electric reported fiscal 2025 results (fiscal year ended February 28, 2026; announced April 2026) of sales around JPY 542 billion (~$3.6B USD) — above its prior

guidance of JPY 525 billion — with operating profit in the area of JPY 48 billion, a high-single-digit operating margin. For context, the prior fiscal year (ended February 2025) had been a down year, with revenue off 6.6% and operating profit down 24.3% — so fiscal 2025 was a recovery year off a weak base. For fiscal 2026 (ending February 2027), Yaskawa guided to sales of JPY 580 billion and operating profit of JPY 60 billion — implying roughly 7% revenue growth and 27% operating-profit growth, i.e., the company expects margin expansion as the cycle improves.

The honest framing on margin: Yaskawa runs a high-single-digit operating margin, materially below Fanuc's low-20s. Yaskawa is the lower-margin of the two Japanese giants — its mix carries more lower-margin systems and standard motion-control product, and it lacks Fanuc's extreme vertical-integration margin structure. That is the trade-off versus Fanuc: Yaskawa is cheaper (34.2x forward P/E versus Fanuc's 46.4x) and arguably better positioned for the humanoid-component angle via its servo dominance, but it is the structurally lower-margin business. Forward consensus among Japanese brokers broadly tracks the company's own guidance for mid-single-digit-plus revenue growth and margin recovery; treat specifics as estimate-basis.

The binary event ahead is the Q1 fiscal-2027 earnings (the quarter ending May 2026), expected around mid-July 2026 — the print that shows whether the fiscal-2026 margin-expansion guidance is on track and whether semiconductor-related motion-control demand (driven by AI investment) is sustaining. With the stock +39.6% above its 50-day MA and RSI at a very stretched 81.7 — the most overbought reading in this entire batch — the recovery-plus-physical-AI trade is not just discovered, it is extended.

Customer mix today

Yaskawa's customer mix is diversified and balanced across its two core segments, with no single-customer concentration risk — a structural positive shared with Fanuc. The Robotics segment (Motoman) sells into automotive and general-industry automation through a global distributor and integrator channel. The Motion Control segment (servos, drives, inverters) sells into a different and partly counter-cyclical set of end-markets: semiconductor manufacturing equipment, machine tools, EV production lines and general factory machinery. Yaskawa does not disclose a top-customer concentration percentage because the structure is genuinely diffuse across thousands of OEM and end-user customers — that diffuseness is real and is a positive.

The structural shift in the mix that matters is segment-level and demand-driven. The fiscal-2025 recovery and the fiscal-2026 guidance lean on motion-control demand tied to semiconductor capital investment — and that semiconductor demand is increasingly AI-

investment-driven, meaning Yaskawa's servo and drive business is becoming a second-order beneficiary of the AI capex wave (AI datacentre buildout drives chip demand, chip demand drives semiconductor-equipment demand, that drives Yaskawa servo and motion-control demand). That is a genuinely favourable mix shift and a structural demand source the pure robot-arm names have less of. On the Robotics side the demand picture is the more standard automotive-and-general-industry automation cycle. The "change is the story" element for Yaskawa: motion control is quietly becoming an AI-capex beneficiary, layered on top of a normal industrial-automation cycle in robots.

What's actually happening in motion control and the AI-capex linkage

The mechanism worth understanding is the indirect AI-capex linkage running through Yaskawa's Motion Control segment, because it differentiates Yaskawa from a pure robot-arm play. The chain: AI datacentre investment drives demand for advanced semiconductors; semiconductor demand drives capital investment in chip-fabrication and chip-assembly equipment; that semiconductor equipment is full of precision servo motors, drives and motion-control systems — a meaningful share of which Yaskawa supplies. Yaskawa management noted that demand in the global semiconductor market has tended to focus on AI-related investment, and that this has been a support for motion-control demand even as some general-industry demand has been mixed. So while Yaskawa is not an "AI stock," its servo franchise has a real, traceable second-order tap into the AI capex cycle.

The robot-side mechanism is the standard automation cycle plus the physical-AI strategic pivot. The near-term robot demand is the automotive-and-general-industry capex cycle, recovering off a weak fiscal-2024. The longer-term mechanism is MOTOMAN NEXT and the SoftBank/NVIDIA partnerships — Yaskawa trying to move its robot business from commoditising programmed-automation hardware up into AI-driven adaptive robots and "physical AI" systems where margin and differentiation are defensible. The honest assessment in May 2026: the AI-capex linkage in motion control is real and is genuinely supporting the print; the physical-AI robot pivot is strategically sound and well-partnered (NVIDIA, SoftBank are serious counterparties) but is still early and not yet a material revenue line. The Q1 fiscal-2027 print is where the durability of the semiconductor/AI-driven motion-control demand gets tested.

The competitive threat / Fanuc, ABB, KUKA and Chinese servo/robot makers

Yaskawa's competitive set has two fronts. On industrial robots, it competes in the global "big four" oligopoly with Fanuc (the higher-margin Japanese rival, covered separately in this batch), ABB and KUKA (Midea-owned) — a relatively stable high-end structure. On Motion Control, Yaskawa competes with a different field: other servo and drive makers including Mitsubishi Electric, Siemens, and — the structurally significant threat — a fast-growing field of Chinese servo-motor makers competing aggressively on price in the world's largest manufacturing market.

The competitive nuance for Yaskawa: it sits in the lower-margin position of the Japanese robot duopoly versus Fanuc, and its servo franchise — while dominant globally — faces the most direct Chinese commoditisation pressure of any business in this batch, because servo motors are a higher-volume, more standardisable product than complete six-axis robots. Chinese servo makers taking domestic share is a real, ongoing margin pressure. The competitive bull case is that Yaskawa's motion-control technology at the high-precision end (semiconductor equipment, advanced robotics) is genuinely hard to commoditise, and the MOTOMAN NEXT / physical-AI move is an attempt to climb above the price war into AI-differentiated systems. There is no material IP litigation defining Yaskawa's competitive landscape as of May 2026. The competitive risk is the same slow grind that pressures Fanuc — Chinese low-end share gain — but Yaskawa is somewhat more exposed to it because servos commoditise faster than complete robots, and that is the main reason Yaskawa's conviction and bucket sit a notch behind Fanuc's quality.

The terminal risk

The terminal risk for Yaskawa is a race between two structural forces: Chinese commoditisation of its servo and robot franchises on one side, and the success of its physical-AI / motion-component pivot on the other. If Chinese servo and robot makers commoditise faster than Yaskawa's MOTOMAN NEXT and physical-AI strategy can lift the mix toward AI-differentiated, higher-margin systems, then Yaskawa's already-modest margin core erodes — and a high-single-digit-margin business with eroding margins is a value trap, not a compounder. This is a more present, more measurable terminal risk than the abstract "humanoids obsolete fixed automation" risk that hangs over Fanuc, because the commoditisation pressure is already visible in the Chinese market today.

The humanoid-disruption terminal risk applies to Yaskawa too, but with an important asymmetry in Yaskawa's favour: because Yaskawa makes the servo "muscle" that goes

inside robots — including humanoids — a world where humanoids succeed is a world that needs vastly more precision actuators and motion control, which is *additive* to Yaskawa's component franchise even as it pressures the Motoman complete-robot business. The SoftBank/NVIDIA physical-AI partnerships and the MOTOMAN NEXT platform are explicitly Yaskawa's roadmap for the transition — a credible one, with serious partners. The named beneficiaries of a "Chinese commoditisation wins" outcome are the Chinese servo and robot champions; the beneficiaries of a "physical AI wins and Yaskawa executes" outcome are Yaskawa itself and NVIDIA. The terminal risk is real and nearer-term than Fanuc's, and it is the binding constraint on the multiple — but the servo-inside-everything position is a genuine structural hedge that a pure robot-arm maker does not have.

Bull / Gap / Optionality

1. The servo-motor dominance is the theme's best structural hedge. Yaskawa makes the precision motion "muscle" that goes inside every robot — industrial, collaborative and humanoid. It does not have to pick the winner of the humanoid race; it supplies the actuators and motion control to whoever wins. That dual position — finished robots plus the muscle inside everyone else's robots — is more theme-resilient than a pure robot-arm franchise.

2. Motion control is a real, traceable AI-capex beneficiary. AI datacentre investment drives semiconductor demand, which drives semiconductor-equipment investment, which is full of Yaskawa servos and drives. Management explicitly cited AI-related semiconductor investment as a support for motion-control demand. Yaskawa is not an AI stock, but it has a genuine second-order tap into the AI capex wave.

3. Cheaper than Fanuc with arguably better humanoid-component leverage. Yaskawa trades at 34.2x forward versus Fanuc's 46.4x, and its servo franchise gives it more direct exposure to the humanoid-actuator supply chain. For an investor who wants the Japanese industrial-robotics-incumbent exposure at a less demanding multiple and with a component-layer angle, Yaskawa is the value-tilted choice.

4. The physical-AI pivot is well-partnered and strategically sound. MOTOMAN NEXT (five industrial robots, two cobots, powered by NVIDIA Isaac and Jetson Orin) plus the November 2025 SoftBank "physical AI at social scale" partnership are serious moves with serious counterparties. Yaskawa is genuinely trying to climb from commoditising programmed automation into AI-differentiated adaptive robotics.

5. Fiscal-2026 guidance is for margin expansion, not just growth. Yaskawa guided to JPY 580 billion sales and JPY 60 billion operating profit for fiscal 2026 — ~7%

revenue growth but ~27% operating-profit growth, i.e., the company expects the cycle recovery to expand margin off a weak base. If delivered, that operating leverage is the near-term re-rating driver.

Gap

1. **The most overbought name in the batch.** RSI 81.7 and +39.6% above the 50-day MA — Yaskawa has run harder than any other name here. The recovery-plus-physical-AI trade is not just discovered, it is extended, and the near-term risk of a sharp mean-reversion pullback is the highest in this batch.

2. **The structurally lower-margin Japanese giant.** Yaskawa runs a high-single-digit operating margin versus Fanuc's low-20s. It lacks Fanuc's extreme vertical-integration margin structure and carries more lower-margin systems mix. The quality gap versus Fanuc is real and is why Yaskawa sits a notch behind on conviction.

3. **Servos face the most direct Chinese commoditisation pressure in the batch.** Servo motors are a higher-volume, more standardisable product than complete six-axis robots, so Chinese servo makers can attack the franchise faster and harder. This is a present, visible margin pressure — not a long-dated abstraction.

4. **The physical-AI pivot is early and not yet a material revenue line.** MOTOMAN NEXT and the SoftBank/NVIDIA partnerships are strategically sound but, as of May 2026, are not yet moving the revenue or margin numbers. The bull case leans on a mix shift that is announced and partnered but not yet proven in the print.

Optionality

Event	Date / window	Direction
Q1 FY2027 earnings (quarter ending May 2026)	~Jul 10, 2026	Binary on margin-expansion guidance + semi/AI demand
AI-driven semiconductor-equipment demand data	Ongoing	Bull — the traceable AI-capex linkage
MOTOMAN NEXT adoption / order momentum	2026	Bull — physical-AI pivot proof
SoftBank physical-AI partnership commercial milestones	2026–2027	Bull if it converts to revenue

Event	Date / window	Direction
Chinese servo-maker share-gain data	Ongoing	Bear — the nearest-term terminal-risk evidence
Humanoid-actuator supply-chain wins for Yaskawa	2026–2028	Bull optionality if disclosed

---

The trade

Yaskawa Electric is the value-tilted, component-leveraged member of the Japanese industrial-robotics pair — cheaper than Fanuc at 34x forward, with a servo-motor franchise that supplies the motion "muscle" inside every robot including the humanoids trying to disrupt it. That servo position is the theme's best structural hedge, and the AI-capex linkage through motion control is a genuine, traceable demand tailwind. The problem is entirely the entry: at RSI 81.7 and +39.6% above its 50-day MA, this is the most overbought name in the batch, and chasing it here is a poor risk/reward. Initiate at JPY 6,656–7,356 (current JPY 7,006 ± 5%) but strongly prefer to leg in — a small starter at most now, with the real position built on a 10%+ mean-reversion pullback that the technical setup makes more likely than not. Size at 1.25% of risk capital — a core-quality sizing but a notch below Fanuc's, reflecting the lower margin structure and the more direct Chinese commoditisation exposure. Stop at JPY 5,950, roughly 15% below current, beneath the structural breakout level — a moderate stop for a quality name where a deeper give-back from a stretched level does not break the durable thesis. The catalyst is the Q1 fiscal-2027 earnings around July 10, 2026, the test of whether the margin-expansion guidance and the AI-driven semiconductor demand are holding. The cleaner expression of the same Japanese-industrial-robotics-incumbent thesis, if you want the higher-quality, higher-margin anchor and are willing to pay the richer multiple, is Fanuc (6954) — covered separately in this batch; own Yaskawa for the value tilt and the servo-component leverage, Fanuc for the margin and balance-sheet quality. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

6954 — Fanuc (6954.T) · BUY (Tier-1) · Conv 7/10 · Bucket B

Fanuc (6954.T)

The world’s largest industrial robot-arm maker — profitable, fortress balance sheet, the incumbent that wins either branch of the robotics future.

Investment Research · Photoncap-style deep dive · v1 of “Fanuc” · May 14, 2026

What Fanuc physically does

Fanuc is the largest industrial robot-arm manufacturer in the world and one of the two or three companies that effectively define factory automation. It makes three interlocking things. First, CNC (computer numerical control) systems — the controllers that run machine tools, where Fanuc has a dominant global share; this is the “brain” of metal-cutting manufacturing. Second, industrial robots — the six-axis articulated arms, painted Fanuc yellow, that weld, paint, handle, assemble and palletise in factories worldwide, from automotive body shops to electronics assembly. Third, “ROBOMACHINE” — Fanuc’s own machine tools (electric injection-moulding machines, wire-EDM, compact machining centres). The three are designed to work together: a Fanuc factory can run Fanuc machines, controlled by Fanuc CNCs, tended by Fanuc robots.

The reason Fanuc is the binding constraint of factory automation: a six-axis industrial robot arm is, for any task that can be brought to a fixed station, faster, more precise, more reliable and cheaper over its life than any humanoid alternative. Fanuc’s arms are the proven, depreciated, bankable workhorses of global manufacturing. The technical moat is reliability and integration — Fanuc robots are sold with extraordinarily long maintenance-free lifespans and a vertically integrated production system where Fanuc’s own robots build Fanuc’s own robots in heavily automated plants in Japan. That vertical integration is why Fanuc sustains operating margins that are the envy of the sector.

The robotics-theme relevance is twofold. Fanuc is the incumbent that the humanoid disruptors are, in the long run, trying to displace — so it is the natural “other side” of the humanoid trade. But it is also a direct beneficiary of the same automation supercycle: Fanuc’s CRX collaborative-robot line extends its reach into the cobot segment and into

small and mid-sized manufacturers, and the EV-supply-chain automation wave is a direct demand driver for its arms. Fanuc does not make humanoids, and that is a deliberate strategic position, not an oversight — covered in the terminal-risk section.

Product roadmap

The industrial-robot line is the core: Fanuc's six-axis articulated arms across the payload spectrum, from small electronics-assembly robots to heavy automotive body-handling arms, plus the SCARA and delta variants for high-speed pick-and-place. These are continuously iterated rather than launched on dramatic cycles — the roadmap here is incremental payload, reach, speed and software improvement.

The roadmap line that matters for the robotics theme is the CRX collaborative-robot series. The CRX range spans payloads from 3kg to 40kg, is designed for intuitive programming and fast redeployment, and is sold with an eight-year maintenance-free positioning — Fanuc bringing its reliability brand into the cobot segment that Universal Robots and Doosan pioneered. The most recent CRX milestone: FANUC America debuted the CRX-3iA, an ultra-lightweight cobot, in April 2026 — a portable, minutes-to-set-up cobot aimed at small, high-mix tasks like welding, screwdriving, inspection and small assembly, expanding Fanuc's addressable market into SMEs that previously could not justify automation. Fanuc has also publicly framed three robotics trends for 2026 around AI-enabled, easier-to-deploy automation.

On the CNC and ROBOMACHINE side, the roadmap is steady iteration of controllers and machine tools. What Fanuc deliberately does *not* have on its roadmap is a humanoid robot. As of May 2026 there is no Fanuc humanoid product or announced humanoid program — Fanuc's stated strategic position is that fixed and collaborative automation, not bipedal humanoids, is where the durable industrial value is. Treat that as a confirmed strategic choice; whether it proves right is the terminal-risk question.

The financial print

Fanuc reported full-year results for the fiscal year ended March 31, 2026 (announced late April 2026): net sales of JPY 857.83 billion (~$5.7B USD), up 7.6% from JPY 797.13 billion the prior year, with net income of JPY 166.54 billion (~$1.1B USD), up from JPY 147.56 billion. For the fiscal year, Fanuc has guided to and historically delivered operating margins in the low-to-mid-20s percent range — for the prior fiscal year ended March 2025, operating income was roughly JPY 158.8 billion on JPY 797 billion of sales,

a ~20% operating margin, and the company targets an industry-leading ~22.5% operating margin supported by near-full automation of its own production lines and vertical component integration. To put that in context: most industrial-automation peers run high-single-digit to mid-teens operating margins; Fanuc's low-20s is structurally superior, and it is backed by one of the strongest balance sheets in Japanese industry — Fanuc carries a very large net cash position and minimal debt.

The growth driver in the FY2026 result was recovery in Europe and EV-supply-chain automation demand; China remains a major and swing-factor end-market. Forward consensus among Japanese and global brokers for the fiscal year ending March 2027 broadly looks for continued mid-single-digit-plus revenue growth as the automation capex cycle improves, though estimates vary; treat any single forward figure as estimate-basis. The forward P/E of 46.4 is rich for an industrial-automation company — it reflects the market pricing Fanuc both for the automation upcycle and for its position as a robotics-theme incumbent, and it is the main valuation caution on the name.

The binary event ahead is the Q1 FY2027 earnings (the quarter ending June 2026), expected around late July 2026 — the print that shows whether the European and EV-automation recovery is sustaining and whether Chinese demand is firming. With the stock +27.5% above its 50-day MA and RSI at a stretched 73.8, the automation-upcycle trade is well-discovered.

Customer mix today

Fanuc's customer mix is one of its defining strengths and the cleanest in this batch — it is genuinely, structurally diversified, with no single-customer concentration risk. The revenue spreads across automotive and EV-supply-chain automation (a major demand bucket — car plants worldwide are core Fanuc customers for body-shop and assembly robots, and the EV transition is driving a fresh automation wave), electronics and semiconductor manufacturing (Fanuc robots and CNCs in chip and device assembly), and general industry (everything from metals to food to logistics). Geographically the mix spans Japan, China, the rest of Asia, the Americas and Europe.

Fanuc does not disclose a top-customer concentration percentage because the structure is genuinely diffuse across thousands of manufacturers and a global distributor/integrator channel — that diffuseness is real and is a positive. The structural shift in the mix to watch is geographic and segment-level rather than customer-level: the FY2026 growth came from European recovery and EV-supply-chain automation, while China — historically a huge Fanuc end-market — has been the cyclical swing factor, soft in parts of the recent cycle and a key variable for the forward print. The CRX

cobot line is also gradually shifting the mix toward smaller end-users — SMEs that the traditional six-axis arm business could not economically reach. The honest framing: Fanuc's customer mix is the picture of a diversified incumbent, and the "change is the story" element here is the cyclical recovery in European and EV automation demand, not a customer-concentration shift.

What's actually happening in the automation capex cycle

The mechanism that drives Fanuc's near-term print is the global factory-automation capex cycle, and the read in May 2026 is mid-cycle recovery. After a soft patch through parts of 2024–2025 — driven by weak Chinese manufacturing demand, EV-investment digestion and general-industry caution — the FY2026 result showed Europe recovering and EV-supply-chain automation re-accelerating. The mechanism is straightforward: when manufacturers commit capex, they buy CNCs, robots and machine tools, and Fanuc, as the scale incumbent with the best margins, captures a large share of that spend. The EV-supply-chain piece is structurally important — building battery plants, motor lines and EV assembly is automation-intensive, and that is a multi-year demand tailwind regardless of the cycle.

The China variable is the swing factor and deserves honest treatment. China is one of Fanuc's largest end-markets, and Chinese manufacturing-capex sentiment has been the single biggest determinant of Fanuc's recent cyclical swings. A firming China is a major upside lever to the forward print; a China that stays soft caps the recovery. Layered underneath the cycle is the CRX cobot mechanism — Fanuc is using the CRX line (and the April 2026 CRX-3iA launch) to pull a new tier of SME customers into automation, expanding the total addressable base rather than just riding the existing one. The honest assessment: Fanuc's near-term story is a well-understood cyclical-recovery-plus-EV-automation story, the company is executing it from a position of structural margin and balance-sheet strength, and the swing variable to watch in the Q1 FY2027 print is China.

The competitive threat / Yaskawa, ABB, KUKA and the Chinese arm makers

Fanuc's direct competitive set is the global industrial-robot oligopoly: Yaskawa Electric (the other Japanese giant — Motoman robots plus dominant servo-motor position, covered as a separate name in this batch), ABB and KUKA (the European players — KUKA

now Chinese-owned via Midea), and a fast-growing field of domestic Chinese industrial-robot makers competing aggressively on price in the world's largest robot market. On the cobot side specifically, Fanuc's CRX competes with Universal Robots and Doosan Robotics.

The competitive picture for Fanuc is the strongest in this batch among the established players. Fanuc, Yaskawa and ABB form a stable oligopoly at the high end where reliability, integration and service matter more than price — and Fanuc's vertical integration and ~20%+ margins give it the best economics of the group. The genuine competitive pressure is from Chinese arm makers taking share in the price-sensitive tier of the Chinese domestic market, where "good enough" robots from local champions are displacing imported ones — this is real and it caps Fanuc's China growth ceiling, but it is a margin/share pressure at the low end, not an existential threat to Fanuc's high-end franchise. There is no material IP litigation defining the competitive landscape for Fanuc as of May 2026. The competitive risk is the slow grind of Chinese low-end share gain, not a sudden displacement — and the high-end oligopoly Fanuc anchors is one of the more durable competitive structures in global industrials.

The terminal risk

The terminal risk for Fanuc is the inverse of the humanoid pure-plays' risk: if general-purpose humanoids actually reach task-economics parity, they structurally erode demand for fixed six-axis arms, and Fanuc's core franchise becomes a slowly declining business. Fanuc's strategic bet — visible in the fact that it has no humanoid program — is that this does not happen on a relevant timeframe: that for the tasks that matter, a fixed arm or a cobot at a station beats a bipedal humanoid on cost and reliability for the foreseeable future. If Fanuc is right, it owns the durable centre of factory automation. If Fanuc is wrong, it is the incumbent being disrupted, and the 46x forward multiple becomes very hard to defend.

The transition timing is the same embodied-AI maturation curve that governs the whole theme. The nuance specific to Fanuc: even in a world where humanoids succeed, they succeed *first* at the tasks fixed automation can't do — operating in unstructured, human-built spaces — which is additive to the automation market rather than purely substitutive, at least initially. The substitution risk to Fanuc's core only bites when humanoids become cheaper and more reliable than fixed arms *at fixed-station tasks*, which is a much higher bar and a much longer-dated risk. The named alternative beneficiaries of a "humanoids win fast" world are the humanoid OEMs (UBTECH, Boston Dynamics/Hyundai, Tesla) and the components supply chain. Fanuc's mitigant is that it could enter the humanoid actuator/controller supply chain if it chose to — it makes precisely the servo and

motion-control technology humanoids need — so even a humanoid-dominant future has a path for Fanuc to participate. The terminal risk is real but long-dated, and it constrains the multiple more than it threatens the near-term earnings.

Bull / Gap / Optionality

1. The structurally most profitable scale incumbent in factory automation. Fanuc runs operating margins in the low-to-mid-20s percent — roughly double the high-single-to-mid-teens of typical automation peers — backed by near-full automation of its own production and a fortress balance sheet with very large net cash. FY2026 net sales of JPY 857.83 billion, up 7.6%, with net income up to JPY 166.54 billion. This is a high-quality industrial compounder, not a speculative theme name.

2. It wins either branch of the robotics future. If humanoids stay uneconomic for fixed-station work, Fanuc owns the durable centre of automation. If the automation supercycle simply continues, Fanuc captures a large share of rising capex. And if humanoids succeed, Fanuc makes the servo and motion-control technology to participate in their supply chain. The asymmetry is favourable across scenarios.

3. EV-supply-chain automation is a structural multi-year tailwind. Building battery plants, motor lines and EV assembly is automation-intensive, and this demand driver runs independent of the broader cycle. FY2026's growth was explicitly driven by EV-supply-chain automation plus European recovery — a tailwind with years left.

4. The CRX cobot line expands the addressable market. The CRX range (3–40kg payload, 8-year maintenance-free) and the April 2026 CRX-3iA ultra-lightweight launch pull a new tier of SME customers into automation — manufacturers the traditional six-axis business could never economically reach. This is genuine TAM expansion layered on the core.

5. A stable high-end oligopoly with durable pricing. Fanuc, Yaskawa and ABB anchor a high-end industrial-robot oligopoly where reliability, integration and service — not price — decide. That competitive structure is one of the more durable in global industrials, and Fanuc has the best economics within it.

Gap

1. A 46x forward P/E is rich for an industrial-automation company. Even a structurally superior compounder is expensive here — the multiple prices both the automation upcycle and a robotics-theme premium. If the cycle disappoints or the theme

premium fades, there is meaningful multiple-compression risk on a stock that is fundamentally a cyclical industrial.

2. China demand is the swing factor and it has been soft. China is one of Fanuc's largest end-markets and the single biggest determinant of its recent cyclical swings. A China that stays soft caps the recovery the current multiple is pricing — and Chinese manufacturing sentiment is hard to underwrite.

3. Chinese arm makers are grinding share at the low end. Domestic Chinese industrial-robot makers compete hard on price in the world's largest robot market, taking the price-sensitive tier. This is a real, ongoing margin-and-share pressure that caps Fanuc's China growth ceiling — a slow grind, but a persistent one.

4. The technically stretched entry. RSI 73.8, +27.5% above the 50-day MA — the automation-upcycle-plus-robotics-incumbent trade is well-discovered and the stock has run hard. Chasing here carries real near-term mean-reversion risk before the next leg.

Optionality

Event	Date / window	Direction
Q1 FY2027 earnings (quarter ending June 2026)	~Jul 28, 2026	Binary on European/EV recovery sustaining + China
China manufacturing-capex recovery data	Ongoing	Bull — the biggest swing lever to the forward print
EV-supply-chain automation order momentum	2026	Bull — structural demand driver
CRX cobot line / CRX-3iA SME-adoption traction	2026	Bull — TAM expansion evidence
Fanuc strategic move into humanoid components	2026–2028	Bull optionality if it comes
Chinese arm-maker share-gain data	Ongoing	Bear — structural low-end pressure

---

The trade

Fanuc is the high-quality anchor of this Asian robotics batch — alongside Hyundai, one of the two names you own without betting the position on whether humanoids work. It is the world's largest industrial robot-arm maker, structurally the most profitable scale player in factory automation, with a fortress balance sheet, and it benefits from the automation supercycle today while sitting on the durable side of the long-term robotics question. The valuation is the catch — 46x forward is rich for a cyclical industrial — so entry discipline matters. Initiate at JPY 7,634–8,438 (current JPY 8,036 ± 5%) and be willing to leg in: the stock is technically stretched (RSI 73.8, +27.5% above the 50-day MA), so take a starter position and add on a 5–10% mean-reversion pullback rather than chasing the full size here. Size at 1.5% of risk capital — a core-quality sizing, justified by the diversified customer base, the balance-sheet strength and the cross-scenario asymmetry. Stop at JPY 6,900, roughly 14% below current, beneath the structural breakout level — a moderate stop appropriate for a high-quality name where the thesis is durable and a deeper give-back does not break it. The catalyst is the Q1 FY2027 earnings around July 28, 2026, with the China end-market read as the key swing variable. If you want a higher-beta expression of the same industrial-robotics-incumbent thesis with more direct humanoid-supply-chain leverage, Yaskawa Electric (6506) — covered separately in this batch, with its dominant servo-motor position feeding the "muscle" of every robot including humanoids and an active SoftBank/NVIDIA physical-AI partnership — is the cleaner shot, with Fanuc the lower-risk, higher-quality anchor. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

SYM — Symbotic Inc. · BUY (Tier-1) · Conv 7/10 · Bucket A

Symbotic Inc. (SYM)

The warehouse-robotics systems integrator that just turned GAAP-profitable on a $22B+ backlog — robotics' best-quality business, with all its risk concentrated in one customer.

Investment Research · Photoncap-style deep dive · v1 of "Symbotic" · May 14, 2026

What Symbotic physically does

Symbotic builds and installs complete warehouse automation systems — not a robot you buy off a shelf, but an entire reconfiguration of how a distribution center physically works. In a traditional grocery or retail DC, pallets arrive, are broken down by hand, products are stored on racks, and human pickers walk miles per shift assembling store-bound orders. Symbotic replaces almost all of that with a system: incoming cases are depalletized (increasingly by Symbotic's own robotic depalletizers), then a swarm of autonomous battery-powered bots — called SymBots — drive at high speed across a dense multi-level steel structure, storing and retrieving individual cases, and assembling outbound pallets that are sequenced and "store-friendly" so they unload in aisle order at the destination store. Software — the part that is genuinely hard — orchestrates thousands of bot movements, inventory positions and order flows in real time.

For the Robotics theme this is the important point: Symbotic is not an early, unprofitable robot OEM hoping a product works someday. It is a systems integrator with a deployed, revenue-generating, now-profitable installed base, and it sits at a real chokepoint — the physical throughput of a retailer's supply chain. That is closer to the "supply chain monetizes first" logic this theme is built on than to the speculative-OEM end. Symbotic's binding-constraint position is the distribution center itself: once a retailer commits a DC to Symbotic, switching is enormously expensive because the building has been rebuilt around the system.

The trade-off baked into the model is fixed infrastructure. Symbotic's system is steel structure plus bots plus software, engineered into a specific building. That delivers

high throughput and deep lock-in, but it is capital-heavy, slow to deploy, and less flexible than wheeled autonomous mobile robots (AMRs) or humanoids that can be added incrementally. That tension is the heart of both the bull and the terminal-risk case.

Product roadmap

Symbotic's core system has scaled through successive generations of the SymBot autonomous case-handling robot and the surrounding structure and software stack, with each generation increasing speed, density and the range of case sizes handled. The product roadmap in FY2026 is defined less by a single new robot and more by two things: vertical integration of the system and expansion into new parts of the warehouse.

The vertical-integration moves are concrete and recent. Symbotic acquired Fox Robotics — a maker of autonomous forklifts — bringing automated trailer loading/unloading into the system rather than leaving it as a manual bookend. And in a landmark transaction, Symbotic agreed to acquire Walmart's Advanced Systems and Robotics (ASR) business, the in-house automation unit Walmart had been developing, simultaneously signing a related long-term commercial agreement. That deal converts a potential in-house competitor into an asset and deepens the Walmart relationship into a multi-decade structural partnership. Symbotic is also developing a new storage/retrieval platform aimed at smaller-footprint and e-commerce fulfillment use cases — the explicit attempt to extend beyond large grocery-style DCs.

The GreenBox roadmap is the warehouse-as-a-service leg: GreenBox is a SoftBank-backed JV in which Symbotic holds a 35% stake, and which has committed to spend at least $7.5 billion on Symbotic systems over roughly six years — a structure that lets end customers buy automation as a service rather than as capex. What Symbotic does not make: standalone humanoids, or generic AMRs sold piecemeal — it sells the integrated system, which is both its moat and its rigidity.

The financial print

Symbotic reported Q1 FY2026 results (the fiscal quarter, reported early February 2026) with revenue of $630 million, up 29% year-on-year, and — the milestone that matters — net income of roughly $13.4 million, the transition to GAAP profitability. Adjusted EBITDA was $67 million, up sharply from $18 million a year earlier, and free cash flow for the quarter was $189 million. Backlog stood at $22.3 billion, the vast majority from Walmart and the GreenBox/Exol-related commitments. Forward guidance from

that print was Q2 revenue of $650-670 million and adjusted EBITDA of $70-75 million, with management signaling sequential acceleration into the back half of the fiscal year.

The most important and most current binary: Symbotic reported its next quarterly results around May 6, 2026. As of this note's price date (May 14, 2026) the stock sits at $48.56 with RSI 34.4 and price 11.0% below its 50-day moving average — a notably weak, oversold tape for a name in a hot theme, which suggests the early-May print or the surrounding commentary disappointed the market or that the Walmart-concentration overhang reasserted itself. That divergence — a profitable, $22B-backlog business trading oversold while the rest of the theme runs — is the single most interesting setup in this batch and is treated directly in the trade section.

Forward P/E is ~63.7 on a ~$29.3 billion market cap. That is a rich multiple, but for a company that just crossed into GAAP profitability with a backlog roughly 8-9x annual revenue and accelerating margins, it is a growth-at-a-real-price multiple, not a pure-narrative one. Sell-side coverage is broad; the bull-bear split centers entirely on Walmart concentration and rollout pace. Next regular earnings after the early-May print: approximately August 5, 2026.

Customer mix today

Symbotic's customer mix is the bull case and the bear case in one number. In FY2025 (ended late September 2025), more than 84% of total revenue came from Walmart — some more recent analysis cites figures around 87%, per Yahoo Finance and Motley Fool coverage in early-to-mid 2026. The next tier: Albertsons, the second-largest US supermarket chain, and C&S Wholesale Grocers, the largest US wholesale grocery distributor, are the named non-Walmart anchor customers, but together they remain a clear minority of revenue. There is no version of Symbotic's 2026 income statement that is not dominated by Walmart.

The structural shift the bull wants is diversification, and the evidence is mixed-but-improving. The GreenBox JV — Symbotic owns 35%, SoftBank-backed, $7.5 billion-plus committed spend over six years — is the primary diversification vehicle, designed to bring Symbotic systems to customers beyond Walmart via a warehouse-as-a-service model. The Exol contract referenced in the Q1 FY2026 results is another non-Walmart data point. And critically, the acquisition of Walmart's ASR business plus the related long-term commercial agreement deepened — rather than diversified — the Walmart relationship, while removing the risk that Walmart builds its own competing system. So the honest read: Walmart is ~84%+ and will remain the dominant customer for years; the diversification is real but slow, and runs primarily through GreenBox.

What’s actually happening at Walmart

Walmart is the entire game, so the mechanism there deserves the detail. Symbotic’s relationship with Walmart spans two programs. The original program is the regional distribution centers — Walmart committed dozens of its US DCs to Symbotic systems, a multi-year buildout that is the source of most of the current $22.3 billion backlog and most of FY2025-2026 revenue. The newer program is e-commerce fulfillment: per the Q1 FY2026 commentary, the backlog tied to the Walmart e-commerce program currently covers only about 400 stores, with management explicitly signaling material upside as commercial rollouts extend beyond the initial prototypes expected within roughly 12 months. That “400 stores, much more to come” framing is the embedded growth option inside the Walmart relationship.

The defining 2026 development is the ASR acquisition. Walmart had been building Advanced Systems and Robotics in-house — a latent threat that Walmart could eventually self-supply automation and stop buying from Symbotic. By acquiring ASR and signing a related long-term commercial agreement, Symbotic neutralized that threat and converted it into capability and contract certainty. That is the most important risk-reducing event in the Symbotic story in years. The flip side: it makes the relationship even more central. The honest read on Walmart as of May 2026 is that the relationship is structurally deeper and more secure than a year ago — but the concentration is, if anything, higher, and the stock’s oversold tape suggests the market is currently weighting the concentration risk and rollout-pace uncertainty over the de-risking.

The competitive threat

Symbotic’s competitive set is the warehouse-automation industry, and it is real but does not directly threaten Symbotic’s installed base because of the switching-cost moat. The named competitors: AutoStore (the Norwegian-listed cube-storage automation specialist, large installed base, different architecture), Ocado (the UK grocery-automation and software player), Dematic (owned by KION) and Honeywell Intelligrated among the traditional material-handling integrators, and a growing field of AMR vendors. AutoStore and Ocado are the closest analogs as pure-play automation businesses; Dematic and Honeywell are the incumbent integrators Symbotic displaces on new projects.

The more strategically interesting competitive vector is architectural rather than corporate. Symbotic’s model is fixed, dense, building-integrated infrastructure. The com-

peting philosophy is flexible, incremental automation — fleets of AMRs (and eventually humanoids) that can be added to an existing warehouse without rebuilding it. For a retailer choosing how to automate a new DC, that is a genuine fork in the road, and some will choose flexibility over throughput. There is no material IP litigation currently driving the thesis. The competitive bottom line: Symbotic's deployed systems are extremely well defended by switching costs, but it must keep winning the next DC against both traditional integrators and the flexible-automation philosophy — and that fight is decided on system economics and rollout track record, both of which Symbotic currently leads on for large-format DCs.

The terminal risk

The terminal risk for Symbotic has two faces. The near-and-concrete face is Walmart: at ~84%+ of revenue, any Walmart decision to slow the rollout pace, defer DC conversions, or — despite the ASR deal — change strategic direction would hit Symbotic's revenue trajectory hard and immediately. The ASR acquisition meaningfully reduced the in-sourcing version of this risk, but it cannot eliminate the pace-and-priority version: Walmart controls the speed at which the backlog converts to revenue.

The longer-tail, structural face is architectural obsolescence. Symbotic's entire value proposition is fixed infrastructure engineered into a building. If warehouse automation evolves decisively toward flexible, incremental systems — large AMR fleets, and eventually general-purpose humanoids that can pick and move cases in a conventional warehouse without the steel structure — then Symbotic's capital-heavy, building-integrated model could become the expensive, rigid option. This is not a 2026 risk; it is a multi-year transition risk, and Symbotic's new smaller-footprint storage platform and the Fox Robotics forklift integration are partial hedges against it. But it is the reason the multiple, even at a now-profitable ~63x forward, is constrained: you are paying for a fixed-infrastructure model in a theme whose loudest narrative is flexible, general-purpose robots. Symbotic has a credible roadmap to extend its system; it does not have a roadmap to become a humanoid company, and it should not try.

Bull / Gap / Optionality

Bull

1. GAAP profitability changes the quality of the name. Q1 FY2026 delivered net income of ~$13.4 million, adjusted EBITDA of $67 million (up from $18 million YoY)

and $189 million of quarterly free cash flow. In a theme full of cash-burning OEMs, Symbotic is a profitable, cash-generative business — this is structurally the highest-quality robotics name in the batch.

2. The $22.3 billion backlog is roughly 8-9x annual revenue. That is years of contracted revenue visibility, the vast majority from Walmart and GreenBox commitments. Few companies in any sector carry a backlog that deep relative to revenue, and it makes the forward growth path far more knowable than the typical robotics story.

3. The ASR acquisition de-risked the existential Walmart threat. By acquiring Walmart's in-house Advanced Systems and Robotics business and signing a related long-term commercial agreement, Symbotic converted its single biggest structural risk — Walmart self-supplying — into capability and contract certainty. This is the most important risk-reducing event in the story in years.

4. The Walmart e-commerce program is a large embedded option. Per Q1 FY2026 commentary, the e-commerce backlog covers only ~400 stores today with management explicitly signaling material expansion beyond the initial prototypes within ~12 months. That is a growth leg inside the existing customer that is not yet in the backlog.

5. The tape is oversold in a hot theme. RSI 34.4 and price 11.0% below the 50-day moving average — for a profitable, $22B-backlog business while the rest of the Robotics theme runs hot. If the early-May print or commentary is the cause and the concentration overhang is the rest, this is the rare name in the theme you can enter without chasing.

Gap

1. Walmart is ~84%+ of revenue — full stop. No diversification narrative changes the near-term arithmetic. Walmart controls the rollout pace, the DC priority, and the conversion of backlog to revenue. Any Walmart deceleration flows straight through Symbotic's income statement, and the ASR deal makes the relationship deeper, not less concentrated.

2. The oversold tape may be the market knowing something. A profitable name trading 11% below its 50-day MA in a momentum theme is a divergence, and divergences are not always opportunities — the early-May print may have flagged a rollout-pace or margin issue the bull case has not fully absorbed. The print needs to be read before sizing up.

3. The fixed-infrastructure model is the rigid option in a flexible-automation world. Symbotic's capital-heavy, building-integrated system is slow to deploy and hard to retrofit. As AMR fleets and eventually humanoids make incremental automation more viable, some retailers will choose flexibility over throughput on their next DC.

4. The multiple still requires execution. At ~63.7x forward earnings, even a now-profitable Symbotic is priced for continued strong execution and margin expansion. A backlog conversion stumble or a margin disappointment has real downside from this level — profitability removes the bankruptcy tail, not the de-rating risk.

Optionality

Event	Date / window	Direction
Q2 FY2026 earnings (reported ~May 6, 2026)	Early May 2026	Binary — likely cause of oversold tape; read before sizing
Q3 FY2026 earnings	~August 5, 2026	Binary on backlog conversion + margin trajectory
Walmart e-commerce rollout beyond 400 stores	~12 months from Q1 FY2026	Bull — large embedded growth option
GreenBox JV deployment milestones	2026-2027	Bull — primary non-Walmart diversification proof
ASR integration progress	2026	Bull if accretive and on-schedule
Non-Walmart customer wins (beyond Albertsons / C&S / Exol)	2026-2027	Bull — directly addresses the concentration overhang

---

The trade

Symbotic is the highest-quality pure robotics business in this batch, and the unusual setup is that you can enter it without chasing — the entry zone is current ±5%, roughly $46.13–$50.99, into an oversold tape (RSI 34.4, price 11% below the 50-day MA) while the rest of the theme runs hot. That divergence is the opportunity and the warning at once: the early-May 2026 print is the most likely cause of the weakness, so the discipline is to read that print's detail before sizing, and to treat the oversold condition as a chance to enter a profitable, $22B-backlog name at a relative discount rather than as a guarantee. Size at up to 2.0% of risk capital — the highest sizing in this batch, justified by GAAP profitability, free cash flow, backlog visibility and the ASR de-risking, but explicitly capped there by the ~84%+ Walmart concentration that no amount of business quality offsets. Stop at roughly $40, below the structural support that has held the

name. The defining catalyst is the just-reported Q2 FY2026 print and then the August 5, 2026 Q3 print — the two windows that show whether backlog is converting to revenue at the guided pace. There is no cleaner expression of the warehouse-automation thesis in this theme; Symbotic is itself the clean expression, and the only reason conviction is not higher is the single-customer gate. **Conviction: 7 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

9880 — UBTECH Robotics (9880.HK) · WATCH (Tier-2) · Conv 6/10 · Bucket C

UBTECH Robotics (9880.HK)

The only listed pure-play that is actually shipping full-size humanoids into car plants — first-mover lead, real revenue, still burning cash.

Investment Research · Photoncap-style deep dive · v1 of "UBTECH Robotics" · May 14, 2026

What UBTECH physically does

UBTECH builds full-size, bipedal, embodied-intelligence humanoid robots — machines roughly human-height and human-shaped that walk on two legs, manipulate objects with two arms and dexterous hands, and run a perception-and-planning AI stack on board. The flagship is the Walker series, currently the Walker S2. The technical proposition is that a general-purpose humanoid can be dropped into a workspace built for humans — a car assembly line, a parts warehouse — and perform tasks (inspection, sorting, part-fitting, material transport) without re-tooling the factory around a fixed robot arm. That is the entire thesis of the humanoid category, and UBTECH is the closest listed company to actually proving or disproving it on a factory floor.

The Walker S2's defining hardware feature is an autonomous hot-swap battery system: the robot can walk to a station, swap its own depleted battery pack, and resume work, removing the single biggest constraint on humanoid uptime — the need for a human to plug it in. This is a unglamorous but commercially decisive detail. A humanoid that runs three hours and then needs a person is a demo; one that self-services its power can theoretically run a shift. UBTECH's software side is a full-stack: the ROSA 2.0 robotic operating framework handles motion and coordination, and the proprietary "Thinker" large model handles task reasoning and natural-language instruction. The company built this stack itself over a decade rather than licensing it, which is why R&D spend has been heavy (RMB 1.9 billion cumulatively 2022–2025) and why losses have been structural rather than one-off.

Beyond the headline humanoid, UBTECH still runs a legacy business: AI education

robots, consumer robots, automated logistics vehicles, and sector-specific service robots. That legacy base was the entire company three years ago; today it is the shrinking remainder of a mix that has flipped decisively toward full-size humanoids. UBTECH does not make its own semiconductors, does not make the harmonic reducers or precision actuators at scale (it sources from the components supply chain), and is not vertically integrated the way Hyundai/Boston Dynamics is becoming. It is a robot OEM and AI-stack developer — a system integrator of the humanoid, monetising the application layer.

Product roadmap

The Walker lineage runs Walker (the original full-size humanoid, shown publicly from 2018), Walker X, then Walker S and Walker S1 (the first industrial-focused generations, deployed for factory pilots through 2024), and now Walker S2 — announced in 2025 with autonomous battery hot-swap, the version that entered mass production and delivery in November 2025. UBTECH stated that since the start of 2025 the Walker series accumulated orders exceeding RMB 1.4 billion (~$195M USD), with the milestone of 1,000 Walker S2 units built marked in early 2026. The company has guided to an annual production capacity target of 5,000 industrial humanoid units in 2026, scaling to 10,000 units in 2027 — and disclosed a partnership framework with Siemens around a 10,000-unit-per-year manufacturing build-out.

Named commercial contracts as of early 2026: a RMB 250 million component-and-solution procurement contract centred on Walker S2, and a RMB 159 million bid win for the Zigong Digital Investment “Humanoid Robot Data Collection Center” — note that this second contract is for a data-collection facility, meaning the robots are being bought partly to generate training data, not purely to do production work. That distinction matters for revenue quality and recurs in the Gap section. UBTECH also disclosed early industrial adoption of Walker S2 by Airbus for aircraft-manufacturing tasks — a Western aerospace reference customer, which is rare for a Chinese humanoid maker.

What UBTECH does not have is a confirmed dexterous-hand generation roadmap with public dates the way it has for the body, and it has not disclosed unit economics per deployed robot beyond the ~RMB 760,000 (~$105k) average selling price implied by 1,079 full-size units sold in 2025. Treat the 2026/2027 capacity numbers as company guidance, not confirmed demand — the orders booked (~RMB 1.4bn) are real, but they are a fraction of a 5,000-unit capacity at ASP.

The financial print

UBTECH closed FY2025 (reported March 25, 2026) with revenue of RMB 2,001.0 million (~$280M USD), up 53.3% year-on-year, and a net loss of RMB 789.8 million (~$110M USD) — a meaningful narrowing from the RMB 1,159.9 million loss in 2024. Gross margin rose to 37.7% from 28.7%, driven by a favourable mix shift toward higher-margin humanoid hardware. The standout line: full-size embodied-intelligence humanoid revenue soared 2,203.7% to RMB 820.6 million, with unit volume up to 1,079 robots; the humanoid segment went from 2.7% of revenue in 2024 to 41.1% in 2025 and is now the company's single largest segment. Capital expenditure rose 53.2% to fund new headquarters and industrial parks in Shenzhen and Wuxi.

For context on the competitive financial profile: rival Unitree posted 2025 revenue of RMB 1.71 billion and an adjusted net profit of RMB 600 million — Unitree is profitable, UBTECH is not. The difference is strategy, not execution failure: Unitree sells mostly sub-$20k research/education units (over 70% of its humanoid volume), while UBTECH sells ~$105k industrial units into automotive plants and absorbs the cost of proving the hardest use case. UBTECH's forward P/E of 67.6 prices in a steep ramp to profitability; consensus modelling is thin and dispersed given the company's stage, so treat any forward multiple as estimate-basis rather than a firm anchor.

The binary event ahead is the H1 2026 interim results, expected around late August 2026 — the print that shows whether the 5,000-unit 2026 capacity guidance is being matched by orders and whether the loss continues to narrow on the mix shift. UBTECH also raised roughly $1 billion in strategic financing through 2025 to fund the expansion, so near-term liquidity is not the risk; sustained cash burn against an unproven demand curve is.

Customer mix today

UBTECH's customer mix has flipped in two years from an education-and-consumer base to an industrial one, and that flip is the story. In 2024, the company's revenue was dominated by AI education robots and consumer/service robots, with full-size humanoids a rounding error at 2.7% of revenue. By FY2025, full-size humanoids were 41.1% of revenue — meaning the customer base is now substantially automotive-manufacturing buyers rather than schools.

The named industrial customers, per company disclosure, are concentrated in Chinese and joint-venture automakers: BYD, Geely Auto, FAW-Volkswagen Qingdao, Audi FAW, Dongfeng Liuzhou Motor, BAIC New Energy, plus logistics deployments at Foxconn

and SF Express. UBTECH has not broken out revenue by individual customer as a percentage — that disclosure does not exist publicly, so any single-customer concentration figure would be an estimate. What is disclosed is the channel concentration: the Walker S series has been deployed at NIO's advanced manufacturing centre and on BYD assembly lines for inspection, assembly and logistics tasks. The Airbus deployment is the one significant non-Chinese, non-automotive reference. The legacy education/consumer/logistics-vehicle segments still exist but are now the minority ~59% of revenue and growing far slower than the humanoid line — the mix continues to shift toward industrial humanoids each reporting period.

What's actually happening at the auto OEMs

The mechanism that matters is whether Chinese automakers are buying Walker S2 units to do real production work or to run pilots and generate training data. The honest read in May 2026 is: a mix of both, weighted toward early-stage. UBTECH's disclosure that Walker S has "proven it can handle inspection, assembly and logistics tasks" on BYD lines and at NIO's manufacturing centre is real deployment language, and the 1,079 units sold in 2025 are booked revenue, not LOIs. But the RMB 159 million Zigong contract is explicitly for a "data collection center" — the robots are being purchased to generate embodied-AI training data, which is a government-and-industrial-policy-driven demand source, not a productivity-ROI-driven one.

The bull mechanism: Chinese automakers are in a brutal price war and are genuinely motivated to cut labour cost, and the Chinese state is pushing humanoid adoption as industrial policy — so UBTECH faces a domestic demand environment that is unusually willing to buy unproven hardware. The "1,000 Walker S2 built" milestone and the move to mass production and delivery from November 2025 are evidence that orders are converting to shipped units. The bear mechanism: a robot bought for a pilot or a data-collection centre does not generate a reorder unless it pays for itself, and UBTECH has not disclosed the per-robot productivity economics that would prove the reorder case. The H1 2026 print is where reorder evidence — or its absence — will start to show. Until then, treat the deployment story as "real pilots at scale, ROI unproven."

The competitive threat / Unitree

The direct competitor is Unitree Robotics, which filed for a Shanghai STAR Market IPO in 2026 targeting roughly $610M at a pre-IPO valuation around RMB 42 billion. The

competitive contrast is stark and goes to the heart of the UBTECH thesis. Unitree posted 2025 revenue of RMB 1.71 billion (up from RMB 392 million in 2024) and an adjusted net profit of RMB 600 million — a 674% profit jump. UBTECH posted higher revenue (RMB 2.0 billion) but a RMB 790 million loss. Unitree gets there by selling mostly low-cost research and education humanoids and quadrupeds — over 70% of its humanoid units went to research/education — while UBTECH sells expensive industrial units into the hardest, lowest-margin-to-prove use case.

The strategic question for an investor is which model wins. Unitree's is capital-light, profitable now, and scales on consumer/research demand; UBTECH's is capital-heavy, loss-making, and bets that the industrial deployment moat — the ROSA stack, the factory references, the Airbus and BYD relationships — becomes durable pricing power once humanoids actually work on a line. Beyond Unitree, the field is crowded: China has over 100 humanoid companies, a number that will consolidate sharply through the IPO wave. UBTECH's defensible claim is that it ranked first globally in full-size humanoid shipments in 2025 — first-mover scale in the specific segment (full-size industrial) that is hardest to enter. There is no active IP litigation of note between the two as of May 2026. The risk is not that UBTECH loses a lawsuit; it is that Unitree's profitable model proves the industrial bet was the wrong one to make.

The terminal risk

The terminal risk for UBTECH is the same one that hangs over the entire humanoid OEM category: the humanoid form factor may never reach task-economics parity with fixed automation for the jobs that actually matter. A six-axis robot arm bolted to a factory floor is cheaper, more reliable, faster and more precise than a bipedal humanoid for any task that can be brought to a fixed station. The humanoid's only structural advantage is operating in a space built for humans without re-tooling — and if that advantage turns out to be worth less than the cost premium and reliability deficit, then UBTECH is a permanently subsidised demonstration company, not a robot OEM with an earnings trajectory.

The transition timing that decides this is the maturation of embodied-AI foundation models — the "Thinker"-class large models — to the point where a humanoid can be given a novel task in natural language and just do it reliably. If that arrives in 2027–2029, UBTECH's installed base, data flywheel and factory references become a real moat. If it slips to the 2030s, UBTECH burns cash for years against capacity it built early. The named alternative beneficiaries of a "humanoids stay uneconomic" outcome are precisely the industrial-arm incumbents — Fanuc, Yaskawa — and the components suppliers who sell into whoever wins. UBTECH does have a credible roadmap (ROSA 2.0,

Thinker, the data-collection deployments are explicitly a data flywheel play), but credibility is not the same as proof, and the multiple you can pay is constrained by the fact that the category's core economic premise is still unproven.

Bull / Gap / Optionality

Bull

1. UBTECH is the listed pure-play with actual humanoid revenue, not a roadmap. FY2025 humanoid revenue of RMB 820.6 million on 1,079 units shipped, up 2,203.7% YoY, with the segment going from 2.7% to 41.1% of mix. No other listed humanoid pure-play has revenue at this scale and trajectory. For a fund building robotics exposure, this is the cleanest listed expression of "humanoids are actually selling" — the optionality is real because the base case is already happening.

2. The loss is narrowing on a genuine mix shift, not financial engineering. Net loss improved from RMB 1,159.9M (2024) to RMB 789.8M (2025) while gross margin rose 900bps to 37.7%, because humanoid hardware carries higher margin than the legacy education base. If the mix shift continues and the 5,000-unit 2026 capacity even half-fills, the path to breakeven is visible rather than hypothetical. The H1 2026 print is the swing factor on this leg.

3. First-mover scale in the hardest segment is a real moat. UBTECH ranked #1 globally in full-size humanoid shipments in 2025. The factory references — BYD, NIO, Geely, FAW-VW, plus Airbus in aerospace — are exactly the relationships a new entrant cannot buy. In a field of 100+ Chinese humanoid startups that will consolidate hard through the IPO wave, being the incumbent at full-size industrial scale is defensible.

4. The demand environment is structurally favourable. Chinese automakers in a price war are motivated buyers of labour-cost reduction, and Chinese industrial policy is actively pushing humanoid adoption (the Zigong data-collection-centre contract is policy-driven demand). UBTECH operates in the one market most willing to buy unproven humanoid hardware at scale — a tailwind a US or European pure-play does not have.

5. Funding is not the near-term risk. UBTECH raised roughly $1 billion in strategic financing through 2025, including Middle East capital. The company can fund the 2026–2027 capacity ramp without a forced raise at a bad price — meaning the bear case has to play out through demand disappointment, not a liquidity crunch, which gives the thesis more time to prove out.

Gap

1. Still loss-making at a 67x forward multiple while a direct competitor prints profit. UBTECH lost RMB 790 million in 2025; Unitree earned RMB 600 million. The market is paying a premium multiple for UBTECH's strategy bet — that industrial deployment becomes durable pricing power — and that bet is unproven. If Unitree's capital-light model is simply the better business, UBTECH re-rates down hard.

2. Order book is a fraction of guided capacity, and order quality is mixed. Cumulative Walker orders since the start of 2025 are ~RMB 1.4 billion against a guided 2026 capacity of 5,000 units (which at ~RMB 760k ASP would be ~RMB 3.8 billion of revenue). The gap is demand that does not yet exist. And part of the booked demand — the RMB 159M Zigong "data collection center" — is robots bought to generate training data, not to do ROI-justified production work. Reorder evidence is the missing proof.

3. No disclosed per-robot unit economics. UBTECH has not published the productivity payback of a deployed Walker S2. Without that, there is no way to underwrite the reorder cycle that turns a pilot fleet into a recurring-revenue OEM. Pilots that do not pay back do not reorder, and the whole bull case rests on reorders.

4. The category's core premise is still unproven. If humanoids never beat fixed automation on task economics for the jobs that matter, UBTECH is a subsidised demo company. This is not a UBTECH-specific flaw — it is the terminal risk of the whole humanoid OEM category — but UBTECH, as the listed pure-play, carries it most directly. The multiple compresses the moment the market doubts the premise.

Optionality

Event	Date / window	Direction
H1 2026 interim results	~Aug 25, 2026	Binary on capacity-fill, loss-narrowing, reorder evidence
Walker S2 5,000-unit 2026 capacity fill-rate disclosure	H2 2026	Bull if orders track capacity
Unitree Shanghai STAR Market IPO	2026	Bear (sets a profitable comp; pressures UBTECH's multiple)
Disclosure of per-robot productivity / payback economics	2026–2027	Bull if economics support reorders

Event	Date / window	Direction
Siemens 10,000-unit/yr manufacturing framework progress	2026–2027	Bull if it converts to capacity + orders
New Western reference customers beyond Airbus	Ongoing	Bull — validates non-China demand

---

The trade

UBTECH is the listed name you own to express “humanoids are actually shipping,” accepting that you are paying a premium multiple for an unproven industrial-deployment bet and a company still losing money. Initiate at HKD 101.65–112.35 (current HKD 107.00 ± 5%, a melt-up-tape entry given RSI is a neutral 48.5 and the stock sits just +1.6% above its 50-day MA — there is no momentum-extension premium to wait out here, which is actually a point in favour of entering now). Size at 1.0% of risk capital — kept deliberately below a full position because the name is pre-profitability, the order book is a fraction of guided capacity, and per-robot economics are undisclosed; this is a Bucket C “chase risk is real, conviction is moderate” sizing. Stop at HKD 90.00, roughly 16% below current and beneath the structural support that has held through 2026 — a wider-than-usual stop because a pre-profit humanoid name is volatile and a tight stop will get shaken out. The catalyst is the H1 2026 interim results around August 25, 2026, where capacity-fill rate, loss trajectory and any reorder evidence will either validate or break the thesis. If you want a cleaner expression of the same humanoid theme without the pre-profit OEM risk, the components supply chain — harmonic reducers, precision actuators, the firms that sell into UBTECH *and* Unitree *and* Boston Dynamics — monetises the theme first and does not carry the binary on whether any single OEM’s strategy is the right one. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

ABB — ABB Ltd · WATCH (Tier-2) · Conv 6/10 · Bucket C

ABB Ltd / ABB Ltd (ABBN.SW / NYSE:ABB) (ABB)

Industrial-automation major handing shareholders a separately listed robotics company in 2026 — the spinoff is the robotics catalyst; everything else is a high-quality electrification compounder.

Investment Research · Photoncap-style deep dive · v1 of "ABB" · May 14, 2026

What ABB physically does

ABB is one of the world's largest electrification and industrial-automation companies — a Swiss-Swedish industrial giant whose products run the unglamorous backbone of modern infrastructure: switchgear, electrical distribution equipment, motors, drives, power-electronics, process-automation systems, and — the reason it sits in this theme — industrial robots. The group reports in three segments: Electrification (the largest, roughly half of revenue — the equipment that distributes and manages electrical power, structurally advantaged by data-center and grid buildout), Motion (motors and drives, roughly a quarter), and Automation (process and discrete automation systems, roughly a quarter). The Robotics division has historically sat within the Automation segment.

For the Robotics theme, the relevant unit is ABB Robotics: the division that makes industrial robot arms — the articulated, multi-axis robots that weld, paint, palletize, assemble and handle materials on automotive lines, in electronics factories, in logistics, and increasingly in general manufacturing. ABB Robotics also makes the surrounding software, controllers and a growing line of autonomous mobile robots and "robotic automation" cells. It is one of the established global leaders in industrial robotics — the Western counterweight, alongside Germany's KUKA (now Midea-owned), to the Japanese duo of FANUC and Yaskawa.

The defining fact for an investor in 2026 is structural, not technical: ABB has announced

it will separate ABB Robotics into its own publicly listed company. So ABB is, right now, in the process of becoming two things — a slimmed industrial-electrification group, and a standalone listed robotics company that current ABB shareholders will own directly. The robotics thesis for ABB is therefore mostly a thesis about that transaction.

Product roadmap

ABB Robotics' product lines span the full industrial-robot range: large six-axis articulated robots (the IRB family) for automotive body-shop and heavy-handling applications; smaller and faster robots for electronics and consumer-goods assembly; delta/parallel robots for high-speed picking; the collaborative GoFa and SWIFTI cobot lines that compete with Universal Robots in the human-adjacent space; and a growing portfolio of autonomous mobile robots following ABB's acquisitions in that area. The software and controller layer — ABB's RobotStudio programming environment and the OmniCore controller platform — is increasingly the differentiator, with the roadmap pushing toward AI-enabled, easier-to-deploy robots and "autonomous" cells that need less specialist programming.

But the roadmap event that actually matters is corporate. ABB announced plans to spin off the Robotics division as a 100%-independent, separately listed company. Per ABB's own announcement, the structure is a share distribution — ABB Ltd shareholders receive shares in the new company (working name "ABB Robotics") as a dividend in-kind, pro-rata to their existing holdings — with likely listing venues Switzerland and Sweden. The originally communicated intention was for the business to begin trading as a separately listed company during Q2 2026. Important nuance and a flag: ABB's own Q1 2026 results commentary (reported April 22, 2026) described the "closing of [the] Robotics divestment" as "expected in the second half of the year." Whether that reflects a modest slip from the Q2 target or simply the difference between "start trading" and "fully close" is not fully clear from public sources as of May 14, 2026 — treat the timing as "2026, around mid-year, with H2 closing" rather than a hard confirmed Q2 date. RBC analysts have suggested ABB Robotics could be worth around $3.5 billion as a standalone, per their 2025 estimate.

What the post-spin ABB stub will not include: the robotics business. Investors buying ABB today for robotics exposure need to understand they are buying the right to receive ABB Robotics shares, after which the ABB stub is a non-robotics electrification/motion company.

The financial print

ABB reported Q1 2026 results on April 22, 2026: group revenue of $8.7 billion versus consensus of roughly $8.4 billion, up 18% year-on-year (up 11% comparable), per Futurum's Q1 FY2026 coverage. Segment detail: Electrification revenue $4.6 billion (+21% YoY), Motion $2.1 billion (+18% YoY), Automation $2.1 billion (+16% YoY) — broad-based strength led by data-center and grid demand. Income from operations was $1.8 billion (20.4% margin), operational EBITA $2.0 billion (23.5% margin), net income $1.3 billion (+20% YoY), and free cash flow $1.3 billion (+92% YoY). The robotics-specific note in the print was simply that the divestment closing is "expected in the second half of the year." Q2 2026 guidance was for comparable revenue growth in the high-single-digit to low-double-digit range with YoY operational EBITA margin improvement.

On the robotics division's standalone economics: ABB has previously disclosed that the Robotics division had 2024 revenues of approximately $2.3 billion — about 7% of ABB Group revenue — with an Operational EBITA margin of around 12.1%, and roughly 7,000 employees and manufacturing hubs in Sweden, the US and China. That ~12% margin is the key number: it is meaningfully below ABB Group's ~23% operational EBITA margin, which tells you industrial robotics is the lower-margin, more cyclical part of the portfolio — part of the strategic logic for separating it.

Forward valuation: ABB trades at roughly 31.4x forward earnings on a market cap of approximately CHF 150 billion, with the reference price CHF 82.86 on the primary ABBN.SW listing (the stock also trades as NYSE:ABB and OTC:ABBNY; the OTC line was quoted around $59.71 in May 2026 coverage — a different share/quote basis, not a price discrepancy on the primary listing). The 1-year return has been strong — RSI 68.6 and price 16.5% above the 50-day moving average mark this as an extended, momentum tape. Next earnings: Q2 2026, expected on or around July 16, 2026 — but the genuine binary for the robotics thesis is the spin-off timing and the standalone valuation ABB Robotics achieves.

Customer mix today

ABB's customer mix at the group level is broad and diversified by design — that diversification is part of why the stock is a quality industrial compounder rather than a volatile single-end-market name. Group customers span utilities and grid operators, data-center developers, automotive manufacturers, process industries (oil, gas, chemicals, mining, pulp and paper), buildings and infrastructure, and discrete manufacturing. No single customer or end market dominates group revenue; the current demand standouts are

data centers and grid investment, which is what powered the Electrification segment's 21% Q1 2026 growth.

The Robotics division specifically — the ~$2.3 billion, ~7%-of-group unit — has historically been weighted toward automotive (body shop, powertrain, increasingly EV-line retooling) plus general industry, electronics and a growing logistics/consumer-goods mix. ABB does not publish a clean per-customer percentage breakout for the robotics division, so this is reporting-basis: automotive has historically been the single largest end market for ABB Robotics, with the strategic effort being to grow the non-automotive "general industry" share to reduce automotive cyclicality.

The structural shift that is "the story" here is not a customer-mix shift — it is the corporate separation. Once ABB Robotics is a standalone listed company, its customer mix becomes a thing investors analyze on its own income statement; inside ABB Group today it is too small a slice to move the consolidated mix. For an ABB shareholder, the customer-mix question that matters is the post-spin stub: a diversified electrification/motion/automation business with data center and grid as the growth engines.

What's actually happening with the spin-off (the event that sets the robotics thesis)

The heart of the ABB robotics story is not a share-gain mechanism at a customer — it is the separation transaction, so that is what this section dissects. The strategic logic ABB has given is clean: ABB Robotics is a ~12% operational-EBITA-margin business inside a ~23%-margin group, it is more cyclical than the rest of the portfolio, and ABB's board believes both companies can grow, attract talent and serve customers better as focused, separately listed entities. The mechanics: a share distribution / dividend-in-kind, so ABB shareholders end up holding both the ABB stub and the new "ABB Robotics" listed entity directly, with Switzerland and Sweden the likely venues.

The timing is the live uncertainty. ABB's original communication targeted the business beginning to trade as a separately listed company during Q2 2026; ABB's April 22, 2026 Q1 results commentary referred to the divestment closing being expected in H2 2026. As of this note (May 14, 2026) the honest read is that the separation is firmly on track for 2026, around mid-year, with full closing in the second half — and an investor should not anchor on a precise day. The valuation question is the other live variable: RBC's 2025 estimate put ABB Robotics around $3.5 billion standalone. Whether the market awards the spun entity a robotics-theme growth multiple or a mature-industrial cyclical multiple is the swing factor — and given the ~12% margin and the FANUC/Yaskawa/KUKA competitive set, a sober base case is "decent industrial multiple, not a humanoid-narrative

multiple.” The honest takeaway: the spin-off is a real, value-surfacing catalyst, but the upside depends on the standalone re-rating, which is not guaranteed.

The competitive threat

ABB Robotics competes in industrial robotics against a well-defined, formidable set. The Japanese duo — FANUC and Yaskawa Electric — are the global volume and margin leaders in industrial robot arms, with FANUC in particular running structurally higher margins than ABB Robotics’ ~12%. KUKA, the German robotics maker now owned by China’s Midea, is the other traditional Western-heritage competitor. And the fastest-moving competitive vector is the wave of Chinese industrial-robot makers (Estun, Inovance and others) taking domestic Chinese share aggressively and beginning to push internationally on price. ABB Robotics is a genuine global leader, but it is a leader in a mature, cyclical, increasingly price-competitive market — not a category creator with pricing power.

There is no material IP litigation driving the thesis. The competitive bottom line is what constrains the spin-off’s re-rating: a standalone ABB Robotics enters public markets as the roughly-fourth-largest industrial-robot player by some measures, with sub-group margins, against FANUC and Yaskawa above it and Chinese entrants pressuring from below. That is a respectable franchise — but it is the competitive reality the market will price, and it argues against assuming the spun entity gets a rich growth multiple just because “robotics” is a hot theme.

The terminal risk

The terminal risk has two layers. For the ABB stub (what most current shareholders will be left holding the most of), the risk is ordinary industrial cyclicality — Electrification, Motion and Automation are riding a genuine data-center and grid-investment wave, but capital-equipment demand is cyclical, and at ~31x forward earnings ABB is priced closer to a structural-growth multiple than a mid-cycle industrial. A capex digestion in data centers or grid would compress both earnings and multiple.

For the spun ABB Robotics entity — the actual robotics exposure — the terminal risk is that industrial robotics is a mature, slow-growth, cyclical, competitive market that does not deserve a growth multiple at all. Industrial robot arms have been sold for decades; the category grows roughly with global manufacturing capex, not exponentially; FANUC and Yaskawa already earn the best margins; Chinese entrants are compressing pricing.

The longer-tail structural risk is that the next leg of factory automation — AI-enabled flexible cells, AMRs, and eventually general-purpose humanoids — shifts value away from the traditional fixed articulated-arm franchise that is the core of ABB Robotics. ABB has cobot (GoFa, SWIFTI), AMR and AI-software roadmaps as partial hedges, but a standalone ABB Robotics is fundamentally a mature-industrial business with an automation-of-the-future option attached. The multiple an investor should pay for the spun entity is constrained by exactly that: a good franchise in a mature market, not a hypergrowth robotics pure-play.

Bull / Gap / Optionality

Bull

1. The spin-off is a real, dated, value-surfacing catalyst. ABB is separating its ~$2.3 billion robotics division into a standalone listed company via a tax-efficient share distribution, on track for 2026. Shareholders end up owning both the focused stub and the robotics entity directly — RBC's 2025 estimate put ABB Robotics around $3.5 billion standalone, value that is currently buried inside a $150 billion conglomerate.

2. The ABB stub is a high-quality compounder riding structural demand. Q1 2026 group revenue of $8.7 billion beat consensus, up 18% YoY, with Electrification (+21%) powered by data-center and grid investment, ~23% operational EBITA margin, and free cash flow up 92% YoY. The non-robotics ABB is a genuinely strong industrial business.

3. Separation should sharpen both businesses. A focused robotics company can pursue robotics-specific M&A, talent and capital allocation without competing for attention inside a giant electrification group; the stub gets a cleaner, higher-margin profile. The strategic logic for the split is sound, not financial engineering.

4. ABB Robotics is a genuine global top-tier player. It is one of the established global leaders in industrial robots, with ~7,000 employees, manufacturing in Sweden, the US and China, and a full product range from heavy six-axis arms to GoFa/SWIFTI cobots and AMRs — a real franchise, not a token unit.

5. Optionality on the standalone re-rating. If the market awards ABB Robotics even a moderate robotics-theme premium versus the ~12%-margin industrial base case, the sum-of-the-parts value exceeds what ABB-combined currently reflects. The spin is a structural upside option with a defined timeline.

Gap

1. The robotics exposure is ~7% of revenue — and then it leaves. ABB Robotics was ~$2.3 billion of group revenue in 2024. An investor buying ABB for robotics gets a small slice, and after the spin the ABB stub has no robotics at all. This is not a robotics holding; it is a temporary claim on a robotics distribution.

2. The spin-off timing has wobbled. The original communication targeted Q2 2026 trading; the April 22, 2026 Q1 results referred to H2 2026 closing. The separation is on track but the precise timeline is soft — anchoring a trade on a hard date is a mistake.

3. The standalone re-rating is not guaranteed — and the base case is sober. ABB Robotics runs ~12% operational EBITA margin versus FANUC and Yaskawa above it and Chinese entrants below. Industrial robotics is mature and cyclical. The market may award the spun entity a perfectly ordinary industrial multiple, in which case the "value-surfacing" upside is modest.

4. ABB itself is on an extended tape at a full multiple. RSI 68.6, price 16.5% above the 50-day moving average, ~31x forward earnings — you would be buying a quality industrial after a strong run, into a price that already embeds the data-center/grid growth and presumably some spin-off optimism.

Optionality

Event	Date / window	Direction
ABB Robotics separate listing begins	~Mid-2026 (originally Q2; H2 closing)	Binary — the core robotics catalyst
ABB Robotics standalone valuation / first trading	At/after listing	Bull if re-rates to a robotics premium; neutral if mature-industrial multiple
Q2 2026 ABB Group earnings	~July 16, 2026	Binary on data-center/grid demand durability
Listing venue / index inclusion detail (CH / SE)	Pre-listing 2026	Binary on liquidity and demand for the spun shares
Data-center / grid capex cycle signals	Through 2026-2027	Bear — the stub's cyclicality risk

Event	Date / window	Direction
Chinese industrial-robot competitive pressure	Ongoing	Bear — caps ABB Robotics' standalone multiple

---

The trade

ABB is a "own it for the spin, then decide" name, and the trade is built around that. Entry zone is current ±5%, roughly CHF 78.72–CHF 87.00 on the primary ABBN.SW listing (NYSE:ABB is the equivalent US line) — with the explicit caveat that the tape is extended, RSI 68.6 and price 16.5% above the 50-day moving average, so you are buying a quality industrial after a strong run, not at a base. Size at 1.25% of risk capital — a moderate position justified by genuine business quality and a dated catalyst, but capped there because the robotics exposure is only ~7% of revenue and leaves the company at the spin. Stop at roughly CHF 70, below the structural support and the 50-day EMA cloud. The defining catalyst is the ABB Robotics separate listing, on track for 2026 around mid-year with H2 closing — the binary is both the timing and the standalone valuation the spun entity achieves. The cleanest way to think about the pivot: ABB is the way to get a free, dated option on a top-tier industrial-robotics pure-play landing on the market, while being paid to wait by a high-quality electrification compounder. But it is not a way to own concentrated robotics exposure — for that, the spun ABB Robotics entity itself, once it trades and once its standalone multiple is visible, will be the cleaner expression, and the right move may be to reassess at that point rather than to treat ABB-combined as the robotics position. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

TER — Teradyne, Inc. · WATCH (Tier-2) · Conv 6/10 · Bucket B

Teradyne, Inc. (TER)

The AI-chip-test leader that also owns the West's largest cobot and AMR franchises — a real robotics business, riding inside a semiconductor story.

Investment Research · Photoncap-style deep dive · v1 of "Teradyne" · May 14, 2026

What Teradyne physically does

Teradyne is two businesses bolted together, and an honest robotics note has to say which one moves the stock. The dominant business — roughly 87% of revenue — is automated test equipment (ATE) for semiconductors: the large, complex machines that chipmakers and outsourced assembly/test houses use to verify that a finished die or packaged chip actually works to spec before it ships. In the AI era this is a structurally advantaged position, because AI accelerators, HBM-stacked memory and the associated complex packages require more test time and more sophisticated test, and Teradyne is the leading Western ATE vendor for exactly that workload.

The second business — the reason Teradyne is in the Robotics theme at all — is Teradyne Robotics, which houses two acquired franchises: Universal Robots (UR), the Danish company that effectively created the collaborative-robot (cobot) category — relatively low-cost, easy-to-program robotic arms designed to work safely alongside humans on factory and SME production lines — and Mobile Industrial Robots (MiR), a maker of autonomous mobile robots (AMRs) that move materials around factories and warehouses without fixed guidance infrastructure. Together UR and MiR are arguably the largest cobot-plus-AMR franchise in the Western world by installed base.

For the theme, the framing must be precise. Teradyne Robotics is a real, sizable, brand-leading robotics business — but it generated $91 million of Teradyne's $1,282 million Q1 2026 revenue, roughly 7%. So Teradyne is a way to own a leading cobot/AMR franchise inside a company whose share price is set by the AI semiconductor-test cycle. The binding-constraint logic that makes Teradyne attractive at all runs through ATE; the robotics segment is the embedded option, not the engine.

Product roadmap

On the robotics side, the product lines are well-defined. Universal Robots' core family is its e-Series and the higher-payload UR20 and UR30 cobot arms, with the broader roadmap pushing toward heavier payloads, longer reach and — the strategically important part — AI-enabled capability, with UR working alongside NVIDIA on AI and "physical AI" tooling to make cobots easier to deploy and more autonomous in unstructured tasks. MiR's line spans small-to-large payload AMRs (the MiR250 through the higher-capacity pallet-moving units) for internal logistics. The roadmap thesis for Teradyne Robotics is that cobots and AMRs are moving from "programmed for one repetitive task" toward "AI-assisted, quickly redeployable," which expands the addressable market beyond large manufacturers into the long tail of smaller factories.

On the ATE side — which is what actually drives the stock — the roadmap is tied to AI compute: test platforms for advanced AI accelerators, for HBM and other advanced-packaging-driven test needs, and for the high-performance compute and networking silicon that the data-center buildout requires. Teradyne's Q4 FY2025 commentary explicitly reframed the company's narrative toward data center and "physical AI" as the twin growth engines.

What Teradyne does not make: humanoids. Universal Robots and MiR are fixed-base cobots and wheeled AMRs respectively — purpose-built, not general-purpose bipedal robots. If the theme's center of gravity moves decisively to humanoids, Teradyne participates only via the cobot/AMR adjacency, not directly.

The financial print

Teradyne reported Q1 2026 results on April 22, 2026: total revenue of $1,282 million, split $1,111 million Semiconductor Test, $91 million Robotics, and $80 million Product Test — with all three segments delivering strong year-on-year growth and AI-driven momentum the explicit driver. The Robotics segment's $91 million was a notable recovery from the weaker quarters of 2025. The context that matters: management's guidance had been extraordinary — per TIKR's 2026 coverage, Teradyne guided Q1 2026 revenue of $1.15-1.25 billion against an analyst estimate of roughly $934 million, a ~34% guide-above-consensus that pointed to a new quarterly record and ~75% YoY growth, driven by the AI test surge. Q2 2026 guidance was again $1.15-1.25 billion.

Forward consensus: per TIKR and public.com coverage in May 2026, the Street carries

roughly 25 buy ratings, 1 hold and 2 sells, with 2026 EPS consensus around $6.26 (a ~58% jump) and a price-target spread from a Street-high near $400 to a Street-low near $250 — the low case reflecting a scenario where AI compute demand digests faster than the multi-quarter boom management implies. The stock traded around $341 on May 4, 2026 and around $360 on May 8, 2026, per Meyka and ad-hoc-news coverage, bracketing the $356.12 reference price used here — the reference price is current and consistent with secondary sources, no staleness flag. Forward P/E is ~37.4 on a ~$55.7 billion market cap. The 1-year return is large — TER has been one of the AI-test winners. Next earnings: Q2 2026, expected on or around July 22, 2026 — the binary, and it is a semiconductor binary far more than a robotics one.

Customer mix today

Teradyne's customer mix has to be read at two levels because the two businesses serve completely different customers. At the group level, revenue is dominated by Semiconductor Test (~87% in Q1 2026), whose customers are the major logic and memory chipmakers and the large outsourced semiconductor assembly and test (OSAT) houses — a concentrated set of very large customers whose capex decisions are tied to the AI compute and advanced-packaging cycle. Teradyne does not break out a clean public per-customer percentage, but the ATE customer base is the familiar roster of leading-edge foundry, fabless-accelerator and HBM-memory players plus the big OSATs.

The Robotics segment — $91 million in Q1 2026, ~7% of group — has the opposite customer profile: a long, fragmented tail of small and mid-sized manufacturers buying Universal Robots cobots and MiR AMRs, often through distributors and integrators, with no single dominant customer. That fragmentation is genuinely a feature — the robotics business has none of the single-customer concentration risk that, say, Symbotic carries with Walmart.

The structural shift worth flagging: through 2025 the Robotics segment was the soft spot, with cobot/AMR demand weak amid a sluggish industrial capex environment; the Q1 2026 recovery to $91 million is the first clear sign of that segment turning. But the dominant mix story for the stock remains the surge in Semiconductor Test on AI demand — that is the change that is "the story," and robotics is along for the ride.

What's actually happening in Semiconductor Test (the segment that sets the stock)

Per the section rules, the heart of the technical edge belongs to the segment that actually moves the name — and for Teradyne in 2026 that is unambiguously Semiconductor Test, not Robotics. The mechanism: AI accelerators and the HBM memory stacked alongside them are vastly more complex to test than conventional chips. Higher transistor counts, advanced packaging (chiplets, 2.5D/3D integration, HBM stacks), and tighter performance binning all increase test time per unit and test sophistication per unit — and test time is ATE revenue. As the industry's leading-edge logic and HBM volumes ramp through the AI buildout, Teradyne's ATE demand ramps with it, which is why management could guide Q1 2026 ~34% above consensus.

The Q4 FY2025 narrative pivot — management explicitly reframing Teradyne around "data center and physical AI" — is the company telling investors that the AI compute cycle, not the broad semiconductor cycle and certainly not robotics, is the demand engine. The risk embedded in that mechanism is the Street-low $250 scenario: AI compute is cyclical, and if accelerator and HBM capacity digestion arrives faster than the "multi-quarter boom" management implies, ATE orders fall as fast as they rose. The honest read as of May 2026: Teradyne is executing exceptionally well into a real AI-test surge, the Q1 2026 print and guidance confirm it, and the open question is durability — how many quarters the boom runs before digestion. The robotics segment, by contrast, is a slow-burning industrial recovery story that will not be the swing factor either way.

The competitive threat

In Semiconductor Test, Teradyne's primary named competitor is Advantest, the Japanese ATE leader — the two effectively form a duopoly at the high end, with Advantest particularly strong in certain memory and SoC test segments and Teradyne strong across logic and a broad platform range. Both are winning in the AI-test surge; the competitive question is share within a growing pie rather than survival. Cohu and others compete in adjacent test and handling niches. There is no material IP litigation driving the thesis.

In Robotics, the competitive set is the one that matters for the theme. Universal Robots competing in cobots faces FANUC, Yaskawa, ABB, KUKA and Doosan Robotics among others — the same industrial-robot majors covered elsewhere in this theme — and a growing field of lower-cost Asian cobot makers. UR pioneered and still leads the cobot category by installed base, but the category has commoditized somewhat and pricing pressure is real. MiR in AMRs competes against a fragmented field of mobile-robot vendors. The competitive bottom line: Teradyne's ATE position is a defensible duopoly riding a strong cycle; its robotics position is category-leading but in an increasingly

competitive, lower-margin space. Neither competitive front is an imminent threat, but the robotics front is the more contested.

The terminal risk

The terminal risk for Teradyne, viewed through the Robotics-theme lens, is almost a definitional one: the robotics segment never becomes the swing factor. At ~7% of revenue, Teradyne Robotics could grow well for years and still not be what determines the stock — TER will continue to trade as an AI-semiconductor-test name. For an investor buying TER as a robotics expression, that is the core risk: you are not really getting a robotics stock, you are getting a semiconductor-test stock with a robotics call option attached, and the option is small.

The terminal risk for the company as a whole is the cyclicality of AI compute test demand. ATE is a capital-equipment business; capital-equipment businesses are cyclical; and the current surge is driven by an AI buildout that will eventually digest. The Street-low $250 price target is the explicit articulation of that — a scenario where compute demand digestion outpaces the boom timeline. There is no clean technology-transition that obsoletes ATE — chips will always need testing — but the demand level is cyclical, and Teradyne is priced (~37x forward) closer to a structural-growth multiple than a cyclical one. The robotics segment's own terminal risk — cobot commoditization and the possibility that general-purpose humanoids eventually encroach on cobot use cases — is real but, given the segment's small revenue weight, not what constrains the multiple. What constrains the multiple is AI-test cyclicality.

Bull / Gap / Optionality

Bull

1. Teradyne is executing into a genuine AI-test surge. Q1 2026 revenue of $1,282 million came after a guide that was ~34% above consensus, with Q2 2026 guided to another $1.15-1.25 billion, per TIKR's 2026 coverage. The AI compute and HBM buildout structurally increases test time and test complexity per unit — this is a real, mechanism-driven demand wave, not a narrative.

2. The ATE position is a defensible duopoly. Teradyne and Advantest split the high-end test market; both are winning in the AI surge. Duopoly economics plus a strong

cycle is a high-quality place to be, and it funds everything else, including the robotics segment.

3. The robotics segment is turning. Teradyne Robotics' $91 million in Q1 2026 was a clear recovery from the weak quarters of 2025. If the industrial capex environment continues to improve and the AI-enabled cobot roadmap (the NVIDIA "physical AI" collaboration) gains traction, the segment becomes a second, lower-cyclicality growth leg.

4. UR and MiR are category-leading, fragmented-customer franchises. Universal Robots created the cobot category and leads it by installed base; MiR is a leading AMR brand. Critically, the robotics customer base is a long fragmented tail — no Symbotic-style single-customer concentration risk anywhere in the segment.

5. The tape is not extended. RSI 50.2 and price only 5.8% above the 50-day moving average — for an AI-cycle winner, that is a neutral, un-stretched setup. You are not chasing a melt-up here.

Gap

1. The robotics segment is ~7% of revenue — you are buying a semiconductor stock. $91 million of $1,282 million in Q1 2026. For an investor whose thesis is robotics, TER simply does not give concentrated exposure to it. The robotics story is real but it is not the story.

2. AI-test demand is cyclical, and the Street knows it. The price-target spread runs from a Street-high near $400 to a Street-low near $250 — the low case being faster-than-expected digestion of AI compute and HBM capacity. At ~37x forward earnings, TER is priced closer to structural growth than to a cyclical, and a digestion quarter would hurt.

3. The cobot category has commoditized. Universal Robots still leads, but faces FANUC, Yaskawa, ABB, KUKA, Doosan and a wave of lower-cost Asian entrants. Pricing pressure in cobots is real, which caps the margin and growth contribution the robotics segment can make.

4. The valuation requires the boom to be durable. Consensus 2026 EPS near $6.26 embeds continued strong AI-test demand for multiple quarters. If the surge proves shorter than the "multi-quarter boom" management implies, both the earnings number and the multiple compress together.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~July 22, 2026	Binary — primarily on AI-test durability, secondarily robotics
Robotics segment quarterly trajectory	Each print	Bull if the $91M recovery sustains and builds
AI compute / HBM capex digestion signals	Through 2026-2027	Bear — the Street-low $250 scenario
UR / NVIDIA “physical AI” cobot roadmap traction	2026-2027	Bull — would re-rate the robotics option
Advantest competitive share shifts	Ongoing	Binary on duopoly share within a growing pie

---

The trade

Teradyne is a high-quality way to own a leading cobot and AMR franchise — but the honest framing is that you are buying an excellent AI-semiconductor-test business with a robotics call option attached, not a robotics pure-play. Entry zone is current ±5%, roughly $338.31–$373.93, and the tape genuinely cooperates: RSI 50.2 and price only 5.8% above the 50-day moving average, so this is one of the un-extended names in the theme. Size at 1.25% of risk capital — a moderate position that reflects real business quality and a reasonable entry, but capped below the theme’s best names because the robotics exposure is only ~7% of revenue and the stock’s fate is set by AI-test cyclicality. Stop at roughly $300, below the structural support and the 50-day EMA cloud. The defining catalyst is the Q2 2026 print around July 22, 2026 — but recognize that it is a semiconductor binary first and a robotics binary a distant second. If the thesis is specifically “cobots and AMRs are a structural growth market,” there is no clean, large-cap pure-play expression of that in Western public equities — which is itself why TER, despite the dilution of the semiconductor exposure, is a defensible way to get the cobot/AMR leadership into the book. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

TSLA — Tesla, Inc. · WATCH (Tier-2) · Conv 6/10 · Bucket C

Tesla, Inc. (TSLA)

The only humanoid program backed by a company that already builds actuators, batteries and inference silicon at automotive volume — priced for a robotics future that has not arrived.

Investment Research · Photoncap-style deep dive · v1 of "Tesla" · May 14, 2026

What Tesla physically does

For the purposes of the Robotics theme, Tesla is not a car company — it is the most vertically integrated humanoid-robot developer in the West, and that is the only lens this note uses. Tesla's car business is covered in Theme 1; here the question is narrow: does Optimus, Tesla's general-purpose bipedal robot, justify any of the roughly $1.66 trillion market capitalization the stock carries, and is the robotics program a real industrial effort or a narrative prop?

Optimus is a 1.7-metre, roughly 57-kilogram electromechanical humanoid. Mechanically it is a stack of three things Tesla already manufactures in volume: electric actuators (rotary and linear, built around the same harmonic and planetary drive principles Tesla uses in EV drive units), lithium-ion cells (a ~2.3 kWh pack derived from Tesla's automotive cell lines), and an inference computer derived from the Full Self-Driving (FSD) compute stack. The Gen 3 design carries 37 joints — nine more than the prior generation — and 22-degrees-of-freedom hands targeting sub-millimetre precision, per Tesla's Q1 2026 commentary. The bet is that the marginal cost of a humanoid is mostly actuators, batteries and silicon, and Tesla is one of very few companies on earth that already makes all three at scale and in-house.

The "binding constraint" framing matters. In the humanoid race the bottleneck is not the chassis — Chinese vendors like Unitree already ship walking hardware cheaply — it is (1) dexterous manipulation hardware that survives industrial duty cycles, and (2) the autonomy software that lets the robot do useful, varied work without teleoperation.

Tesla's claimed edge is that its FSD vision-and-neural-net stack transfers to manipulation, and that its factories give it a captive, instrumented training environment. Whether that edge is real is the entire debate. As of May 2026 Optimus generates essentially zero external revenue; it is a pre-product. The stock therefore prices Optimus as an option, not a business.

Product roadmap

Optimus has moved through Gen 1 (the 2022 "Bumblebee" prototype and the 2022 AI Day reveal), Gen 2 (demonstrated December 2023, improved hands and balance), and is now mid-transition to Gen 3 — described by Tesla as "the first design meant for mass production." The Gen 3 reveal has slipped repeatedly: originally targeted for Q1 2026, on the April 22, 2026 Q1 earnings call Elon Musk pushed the reveal to "probably middle of this year," per Electrek's April 22, 2026 coverage. That slip is itself a data point — it is the second consecutive deferral of the Gen 3 unveiling.

The manufacturing roadmap, as stated by management on the Q1 2026 call: Optimus production begins at Fremont in "late July or August" 2026, on the line being vacated as Model S and Model X end production in early May 2026. The stated capacity target is a 1-million-unit-per-year line, with low-volume production starting before end-2026 and a "major scale-up" through 2027. A second dedicated facility at Giga Texas is under construction, with production targeted for summer 2027 and a longer-run ambition of up to 10 million units per year, per The Robot Report's 2026 coverage. Treat the 1M and 10M figures as aspiration, not guidance — Tesla has no track record of hitting humanoid timelines, and the Gen 3 slips undercut the near-dated ones.

What Tesla does not yet make: a commercially sold, externally deployed humanoid. Every Optimus unit through 2026 is destined for internal Tesla factory use ("the robot builds the robot" framing). External commercial sales are a 2027-and-beyond story by Tesla's own account.

The financial print

Tesla reported Q1 2026 results on April 22, 2026: total revenue of $22.38 billion, up 16% year-on-year from $19.3 billion, with automotive revenue of $16.2 billion (also +16%) and energy revenue of $2.41 billion (down 12%). Adjusted EPS was $0.41 versus $0.37 expected — a modest beat — and gross margin expanded 478 basis points YoY to 21.1%, per CNBC and Electrek's April 22, 2026 reporting. Q1 deliveries were 358,023 vehicles,

missing the consensus by roughly 7,600 units. Capex guidance for 2026 was raised to "over $25 billion" from a $20 billion prior figure — a 67% jump in quarterly capex — funding robotaxi, AI compute and Optimus tooling.

None of that print is robotics revenue. Optimus contributes nothing to the $22.38 billion. The robotics line in any Tesla model is pure forward assumption. Forward consensus: the stock trades at roughly 175.7x forward earnings on a market cap of ~$1.66 trillion, a multiple that is uninterpretable on the car business alone (peers trade single-digit to low-20s) and is only rationalizable if the market is capitalizing robotaxi and Optimus as large future businesses. Sell-side is split: bulls such as Morgan Stanley have historically ascribed tens-to-hundreds of dollars per share to Optimus in sum-of-the-parts models; skeptics treat it as a zero. The 1-year return into May 2026 has the stock up meaningfully but well off intra-year highs.

Next earnings: Q2 2026 drops approximately July 22, 2026 — but the more important binary for the robotics thesis is the Gen 3 reveal and the Fremont line start, both targeted for July/August 2026.

Customer mix today

Optimus has no external customer mix — this is the honest answer, and the section is short because the reality is thin. The first and only "customer" through 2026 is Tesla itself: Optimus units will be deployed inside Tesla's own factories doing material handling and simple assembly tasks, by Tesla's own account on the Q1 2026 call. There is no third-party revenue, no signed external purchase order disclosed, and no pilot-customer percentage to break out.

This contrasts sharply with the warehouse-robotics names in this theme (Symbotic has Walmart at ~84%+ of revenue; Teradyne's robotics arm sells cobots to thousands of SMEs). Tesla's robotics "customer concentration" is 100% internal. The structural shift to watch — the thing that would actually validate the thesis — is the first disclosed external order, which Tesla guides to 2027+. Until that prints, the customer-mix story is a promise.

What's actually happening inside Tesla's factories

The mechanism that matters for Optimus is the captive-deployment flywheel: Tesla intends to use its own production lines as both the first market and the training ground.

The logic is that a factory is a structured, repetitive, heavily instrumented environment — the easiest possible place to get a humanoid doing useful work, and the place where Tesla can collect manipulation training data at scale without paying a customer to tolerate failures.

The concrete near-term events: Model S/X production ended at Fremont in early May 2026, freeing the line; Optimus production is guided to start on that line in late July or August 2026. The first units are for internal use. The credibility test is simple and dated — by end-2026 we should see Optimus units actually working on Tesla lines in disclosed numbers, not in choreographed demos. Tesla has a documented history of demonstrating Optimus in teleoperated or heavily-staged settings (the October 2024 "We, Robot" event drew criticism that the robots were remote-operated). The bull needs autonomous, repetitive factory work shown with unit counts; the bear points to the Gen 3 reveal slipping twice as evidence the hardware is not production-ready. As of May 2026 the honest read is: unproven, with the proof point only weeks away.

The competitive threat

The humanoid field in 2026 is genuinely crowded, and Tesla is not the clear leader on hardware. Boston Dynamics (owned by Hyundai) unveiled a production electric Atlas at CES 2026 — 56 degrees of freedom, 50 kg lift, autonomous battery swap — and per CES 2026 coverage every 2026 unit is already committed, shipping to Hyundai's Metaplant and to Google DeepMind. Figure AI's Figure 03, running its Helix 02 autonomy stack, is demonstrating continuous unsupervised operation and has a 2026 roadmap of factory deployments. Agility Robotics' Digit is in active production with 2026 units committed and Hyundai deployments targeted for 2028. On the cost end, China's Unitree shipped 5,500+ units in 2025 and targets 10,000-20,000 in 2026, and UBTech is scaling fast — per the Yahoo Finance "Humanoids Market" report, Unitree, UBTech and Boston Dynamics are characterized as current leaders with Tesla and Agility as scaling challengers.

The honest competitive read: Tesla is behind Figure and Boston Dynamics on demonstrated autonomy and behind Unitree on shipped volume and cost. Tesla's structural advantage is not first-mover — it is manufacturing scale and vertical integration. If humanoids become a volume game decided by cost-per-actuator and in-house cell supply, Tesla wins. If they become an autonomy game decided in the next 24 months, Figure and Boston Dynamics are ahead today. There is no active IP litigation that is material to the thesis.

The terminal risk

The terminal risk for Optimus is not a technology transition that obsoletes it — it is that the program never reaches a working product with viable unit economics, and the option the market has capitalized goes to zero. Tesla stock at ~175x forward earnings is not priced on cars; a meaningful slice of the $1.66T cap is the market paying today for a robotaxi-plus-Optimus future. If Optimus stalls — if Gen 3 keeps slipping, if the factory deployment is quietly teleoperated, if the autonomy does not generalize beyond demos — that slice deflates.

The longer-tail structural risk is that the humanoid form factor itself proves to be the wrong abstraction for most industrial work. Purpose-built automation (Symbotic's warehouse systems, fixed-base cobots from Universal Robots, wheeled AMRs) is cheaper, more reliable and already deployed. The humanoid bet is that general-purpose beats purpose-built because it slots into human-designed environments without retooling. That may be true for some tasks and false for most. Tesla has a credible roadmap on hardware cost; it does not yet have proof that general-purpose autonomy is solved. The multiple you can pay is constrained by the fact that the entire robotics leg is, today, faith-based.

Bull / Gap / Optionality

Bull

1. Vertical integration is a real, structural cost advantage. Optimus's bill of materials is dominated by actuators, battery cells and inference silicon — all three of which Tesla manufactures in-house at automotive volume. No other Western humanoid developer has captive cell supply or its own AI training/inference silicon. If humanoids become a cost-driven volume market, Tesla's marginal-cost position is the best in the West. This is the leg of the thesis with the firmest foundation.

2. The captive factory deployment de-risks the go-to-market. Tesla does not need to win an external sales cycle to deploy its first thousands of units — it deploys into its own factories, which it controls, and which double as a training-data engine. The Fremont line start in July/August 2026 is the first concrete proof point and it is weeks away, not years.

3. Manufacturing scale is exactly the capability the field lacks. Figure, Agility and Boston Dynamics are excellent at robots and unproven at mass manufacturing. Tesla has built millions of complex electromechanical machines. The 1M-unit-line ambition is

aspirational, but the underlying claim — that Tesla can industrialize a humanoid faster than autonomy-first startups — is plausible and is the source of the option value.

4. The FSD stack is a genuine autonomy asset. Tesla's vision-based neural-net inference stack, whatever one thinks of robotaxi timelines, is a large, real, deployed AI system. If even part of it transfers to manipulation and navigation, Tesla starts the autonomy race with a running engine rather than from zero. Per Tesla's Q1 2026 commentary the company is explicitly architecting Optimus autonomy off the FSD foundation.

5. The capital base removes funding risk. Tesla raised 2026 capex guidance to over $25 billion and runs a profitable core business. Optimus can be funded through years of losses without dilution or a financing cliff — an advantage no pure-play humanoid startup has.

Gap

1. The Gen 3 reveal has slipped twice. Originally Q1 2026, then pushed to "middle of this year" on the April 22, 2026 call, per Electrek. For a program management insists is near mass production, repeatedly deferring the reveal of "the first design meant for mass production" is a credibility problem. Schedule slips at the prototype stage usually predict schedule slips at the production stage.

2. The valuation gives zero margin of safety. At ~175.7x forward earnings and a $1.66T cap, the stock already capitalizes a successful Optimus and a successful robotaxi. A clean factory-deployment proof point in 2026 may be necessary just to hold the stock, not to re-rate it. The asymmetry is poor: the good case is largely priced, the bad case is not.

3. Demonstrated autonomy lags Figure and Boston Dynamics. Tesla's public Optimus demonstrations have repeatedly been criticized as teleoperated or staged (the October 2024 "We, Robot" event most notably). Competitors are showing continuous unsupervised operation today. On the metric that actually decides the humanoid race in the near term, Tesla is not in front.

4. The core business is decelerating. Q1 2026 deliveries missed by ~7,600 units and energy revenue fell 12%. The auto growth story is, per Electrek's own framing, "dead." That matters because it removes the fallback: if Optimus disappoints, there is no re-accelerating car business to catch the stock.

Optionality

Event	Date / window	Direction
Optimus Gen 3 reveal	“Middle of 2026” (slipped from Q1)	Bull if specs/dexterity impress; Bear if slips again
Fremont Optimus line start	Late July / August 2026	Binary — first real production proof point
Q2 2026 earnings	~July 22, 2026	Binary on robotics commentary + capex detail
First disclosed Optimus factory deployment counts	H2 2026	Bull if autonomous + numbered; Bear if vague/teleoperated
First external Optimus order	2027+	Bull — validates the customer-mix thesis
Giga Texas Optimus facility production start	Summer 2027	Bull on scale-up credibility

---

The trade

Tesla as a robotics expression is a hold-and-watch, not a chase. The entry zone is current ±5%, roughly $419.84–$464.04, but with the explicit caveat that there is no valuation margin of safety here — you are buying an option on Optimus inside a stock that already prices the option richly. Size at 1.5% of risk capital, treating this as the diversified large-cap leg of the Robotics theme rather than a high-conviction single-name bet; the genuinely asymmetric robotics exposure in this theme sits in the supply chain, not in the OEM at 175x. Stop at roughly $360, below the recent structural support and the 50-day moving-average cloud. The defining near-term catalyst is the cluster in July/August 2026 — the Gen 3 reveal and the Fremont line start — which together are the first hard test of whether Optimus is an industrial program or a narrative. If the thesis is “humanoids are real and the supply chain monetizes first,” the cleaner expression is the actuator/component names in the theme’s supply-chain bucket, which monetize Tesla’s Optimus build regardless of whether Tesla’s own autonomy works. Tesla is the way to own the OEM optionality with funding risk removed; it is not the way to own robotics cheaply. **Conviction: 6 / 10.**

Sources referenced inline throughout. Reference v1 of this template format:

_Watchlist/hanmi-photoncap-style.md.

277810 — Rainbow Robotics (277810.KQ) · WATCH (Tier-2) · Conv 5/10 · Bucket C

Rainbow Robotics (277810.KQ)

Samsung’s in-house humanoid and cobot developer — strategic-owner optionality, priced as if the optionality is already cashed.

Investment Research · Photoncap-style deep dive · v1 of “Rainbow Robotics” · May 14, 2026

What Rainbow Robotics physically does

Rainbow Robotics is a Korean robotics developer with a KAIST pedigree — it was founded in 2011 by the research team out of KAIST’s Humanoid Robot Research Center, the group that built “Hubo,” the first Korean bipedal walking robot. That heritage matters because it is the genuine humanoid IP foundation underneath what is, today, mostly a collaborative-robot and service-robot business. Rainbow makes three things: collaborative robot arms (the RB cobot series), service and mobility robots (including coffee-making robots and quadrupeds), and — the reason the stock trades where it does — humanoid and semi-humanoid platforms descended from the Hubo lineage.

The technical centrepiece in 2026 is the RB-Y1, a wheeled semi-humanoid. It is not a bipedal walker; it is a torso with two seven-degrees-of-freedom arms, a six-DoF “single leg,” mounted on a wheeled mobile base — a mobile dual-arm manipulator. The design choice is deliberate: a wheeled base is far more stable, energy-efficient and reliable than bipedal legs for the indoor manipulation tasks (logistics, manufacturing assist, lab automation) that are commercially reachable today, while still giving the robot a human-like reach-and-manipulate envelope. Rainbow has also developed omnidirectional wheels and a development kit for the platform, positioning RB-Y1 as a research-and-development platform that customers — most importantly Samsung — buy to build applications on, rather than a finished product sold for a fixed job.

The strategic fact that dominates everything else: Samsung Electronics owns roughly 35% of Rainbow Robotics and is its largest shareholder, having lifted its stake from a

minority position to controlling-influence level through 2024 and into 2025, and established a “Future Robotics Office” to direct the relationship. Rainbow is, functionally, Samsung’s in-house robotics arm with a separate listing. It does not make its own chips or AI foundation models — the Samsung relationship is explicitly about pairing Samsung’s AI/software with Rainbow’s robot hardware. Rainbow is a hardware-and-controls company; the intelligence layer is the Samsung side of the partnership.

Product roadmap

The cobot line is the RB series — RB collaborative arms across a payload range, the legacy revenue base, sold into Korean and international SMEs for machine-tending, assembly and packaging. The service-robot line includes the coffee-robot and quadruped products that built early revenue. The roadmap event that re-rated the stock is the humanoid push: the RB-Y1 wheeled semi-humanoid was rolled out in April 2025, and an enhanced, faster, more agile RB-Y1 research platform — with omnidirectional wheels and a development kit — followed through 2025. Samsung has signalled it will deploy collaborative variants of RB-Y1, plus Rainbow’s RB cobots, dual-arm mobile manipulators and autonomous mobile robots (AMRs), inside its own manufacturing and logistics operations.

The genuinely confirmed items are: the 35% Samsung stake, the Future Robotics Office, the RB-Y1 platform and its enhancements, and Samsung’s stated intent to deploy Rainbow robots in its own facilities. What is not confirmed with public dates is a bipedal humanoid commercial product — Rainbow has the Hubo bipedal heritage and the engineering capability, but as of May 2026 there is no announced bipedal humanoid product with a launch date and a price. The RB-Y1 is the real near-term product, and it is being sold as a research/development platform. The pathway from “Samsung deploys RB-Y1 collaborative variants in its fabs” to “Rainbow books recurring high-margin humanoid revenue” is the roadmap’s missing middle, and it is undated. Treat the deployment intent as confirmed and the revenue conversion as estimate-basis.

The financial print

Rainbow Robotics reported FY2025 revenue of KRW 34.12 billion (~$24M USD), up 76.4% from KRW 19.35 billion in 2024 — strong top-line growth off a tiny base. But FY2025 net profit was approximately KRW 1.42 billion (~$1M USD), down 33.4% year-on-year: the company grew revenue fast and earnings fell, because the humanoid R&D

build-out and the cost of scaling are running ahead of profit. This is the central tension of the financial print — Rainbow is a real, growing, technically credible robotics company, but it is a ~$24M-revenue, barely-profitable one carrying a market capitalisation of KRW 16.3 trillion (~$11.5B USD). The forward P/E of 707.9 is not a typo; it is what happens when a near-zero-earnings company is valued for a strategic-owner humanoid option.

For comparison, the company's valuation has been pegged in the $5.8B–$11.5B range over recent quarters depending on the share price, almost entirely a function of the Samsung relationship and humanoid-theme sentiment rather than the printed financials. Consensus modelling at this stage is essentially a sentiment exercise — there is no robust multi-broker FY2026 earnings consensus to anchor to, because the earnings are too small and too dependent on undated Samsung deployment economics. Treat any forward number as estimate-basis.

The binary event is the Q1 2026 earnings, due around May 15, 2026 — the day after this note's price date. That print will show whether revenue growth is holding and, critically, whether any Samsung deployment revenue is starting to appear in the numbers or whether the stock is still being held up purely by the ownership story. The stock's technical setup amplifies the binary: RSI 71.6 and +26.6% above its 50-day MA going into the print means the easy money has been made and the bar is high.

Customer mix today

Rainbow Robotics' customer mix is the inverse of what its valuation implies. The valuation says "Samsung's humanoid champion"; the revenue says "small Korean cobot and service-robot maker." The legacy revenue base — the ~KRW 34 billion of FY2025 sales — comes from RB cobot arms sold to Korean and international SMEs, plus service robots (coffee robots, quadrupeds) and AMRs. Rainbow does not break out revenue by named customer as a public percentage, so a precise customer split would be an estimate; the structural fact is that the *printed* revenue is broad-based small-ticket industrial and service sales, not Samsung.

The structural shift — and the entire investment story — is Samsung moving from outside investor to 35% controlling shareholder and stating it will deploy Rainbow's RB-Y1 collaborative variants, RB cobots, dual-arm manipulators and AMRs inside Samsung's own manufacturing and logistics operations. If that intent converts, Samsung becomes the dominant customer by revenue, transforming the mix from "diversified SME base" to "captive single-customer with a very long runway." As of the most recent disclosures in early 2026, that conversion has not visibly shown up in the printed numbers at scale — the RB-Y1 is still being sold predominantly as a research platform, and Samsung's

deployment is at the intent-and-pilot stage rather than the volume-procurement stage. The mix today is legacy; the mix the valuation prices is future and undated.

What’s actually happening at Samsung

Samsung is the only customer relationship that matters for the Rainbow thesis, and the honest mechanism is this: Samsung has made a strategic ownership commitment but has not yet made a disclosed volume-procurement commitment. The confirmed facts — Samsung at ~35%, the Future Robotics Office established to direct robotics strategy, and Samsung’s public statements that it will use RB-Y1 collaborative variants, RB cobots, dual-arm mobile manipulators and AMRs to automate its manufacturing and logistics — are real and they are why the stock is a humanoid name rather than a small-cap cobot stock. Samsung’s motivation is also real: it is racing Hyundai/Boston Dynamics, Tesla, and the Chinese field, and it needs a hardware platform; Rainbow’s KAIST/Hubo heritage gives it credible bipedal IP without building from scratch.

But the mechanism that turns this into Rainbow shareholder value is volume procurement at a margin, and that is the part with no disclosed tool counts, no deployment-unit numbers, and no dated procurement schedule. Samsung deploying RB-Y1 collaborative variants “inside its fabs” could mean a handful of pilot units or a multi-thousand-unit fleet — the disclosure does not distinguish. The bull mechanism is that a 35% owner has every incentive to scale its own captive supplier and Samsung’s manufacturing footprint is enormous. The bear mechanism is that Samsung, as a 35% owner, also captures most of the upside of cheap robots through its *own* P&L, not through Rainbow’s — and a strategic owner can keep a captive supplier sub-scale and low-margin indefinitely if that serves Samsung. The Q1 2026 print on May 15 is the first place to look for actual Samsung deployment revenue. Until it appears, “what’s happening at Samsung” is intent, not numbers.

The competitive threat / Doosan Robotics and the global cobot field

Rainbow’s competitive set is two-tier. On the cobot side — its actual revenue base — it competes with Doosan Robotics (the larger Korean cobot maker, also pivoting to humanoids), Universal Robots (the global cobot leader), and Fanuc’s CRX line. On the humanoid side — its valuation base — it competes with the entire humanoid field: Hyundai’s Boston Dynamics (Korean, vertically integrated, far ahead on bipedal Atlas),

UBTECH and Unitree in China, Tesla Optimus, and the US startup field. Rainbow is sub-scale on the cobot side relative to Doosan and Universal Robots, and behind Boston Dynamics on the humanoid side.

The competitive argument for Rainbow is not that it out-engineers the field — it is that it is the one with Samsung's balance sheet, AI stack and captive manufacturing demand behind it. That is a genuine differentiator: a humanoid developer's biggest problems are funding and a first customer, and Rainbow has solved both via the 35% owner. But it is also a competitive *limitation* — Rainbow's ceiling is partly set by Samsung's strategic decisions, where Boston Dynamics under Hyundai has an owner that is building a 30,000-unit/year factory and committing its entire 2026 Atlas production to internal deployment. Both Korean humanoid efforts are owner-captive; Hyundai's is moving faster and committing harder. There is no active IP litigation of note involving Rainbow as of May 2026. The competitive risk is not a lawsuit — it is being a sub-scale cobot maker on printed revenue while priced as a humanoid leader, in a field where the actual humanoid leaders are pulling ahead.

The terminal risk

The terminal risk for Rainbow Robotics is owner-dependency cutting the wrong way: Samsung concludes, after a few years of pilots, that the humanoid bet does not pay on its timeline — or that buying robots from a more advanced supplier is better than scaling its captive one — and slows the funding and deployment commitment. In that scenario Rainbow reverts to what its printed financials say it is: a sub-scale, low-margin Korean cobot and service-robot maker, and the KRW 16 trillion valuation has nothing underneath it. The stock does not need a technology to be obsoleted; it needs Samsung's enthusiasm to cool.

The transition timing that governs this is the broader embodied-AI maturation curve — the same one that governs UBTECH and every humanoid name. If general-purpose humanoids reach real task-economics parity in 2027–2029, Samsung will scale Rainbow hard and the valuation is justified in hindsight. If parity slips to the 2030s, Samsung — a company managing memory-cycle headwinds and capital discipline — has every reason to throttle a speculative captive R&D programme, and Rainbow is left stranded. The named alternative beneficiaries of a "Samsung slows down" outcome are the more advanced humanoid players (Boston Dynamics) and, as always in this theme, the components supply chain that sells into whoever survives. Rainbow does have credible engineering and the Hubo heritage — it is not a fake company — but its multiple is underwritten by a single owner's continued strategic conviction, and that is the most fragile kind of multiple support.

Bull / Gap / Optionality

1. Samsung at 35% is a genuine, rare strategic anchor. A humanoid developer's two hardest problems are funding and a first customer; Rainbow has solved both. Samsung is the largest shareholder, has set up a Future Robotics Office to drive the relationship, and has publicly committed to deploying Rainbow's robots in its own manufacturing and logistics. No other listed Asian humanoid pure-play has this calibre of owner backing it.

2. Revenue is growing fast — 76.4% in FY2025. FY2025 revenue of KRW 34.12 billion was up from KRW 19.35 billion, evidence the cobot/service base is real and accelerating even before Samsung deployment revenue shows up. The growth rate, if sustained, eventually grows into a less absurd multiple — though it has a very long way to go.

3. The RB-Y1 is a credible, shipping product, not a render. The wheeled semi-humanoid rolled out April 2025, was enhanced through 2025 with omnidirectional wheels and a development kit, and is being sold as a research/development platform. The wheeled-base design choice is the commercially sensible one for near-term indoor manipulation — Rainbow is not betting everything on unsolved bipedal locomotion.

4. KAIST/Hubo heritage is real bipedal IP. Rainbow built the first Korean bipedal walking robot. If and when Samsung wants to push from wheeled semi-humanoid to full bipedal, the foundational engineering exists in-house — it does not have to be acquired or licensed. That is genuine optionality on top of the RB-Y1 base.

5. Samsung's manufacturing footprint is the demand runway. If Samsung converts intent to volume procurement, the deployment runway across Samsung's global fabs and logistics is enormous and captive. The bull case does not need Rainbow to win the open market — it only needs Samsung to scale its own captive supplier, which a 35% owner is structurally incentivised to do.

Gap

1. 707x forward P/E for a $24M-revenue, barely-profitable company. FY2025 net profit was ~KRW 1.42 billion, *down* 33.4% YoY, against a KRW 16.3 trillion market cap. The valuation is not supported by any version of the printed financials — it is a pure strategic-option premium, and option premiums compress fast when sentiment turns.

2. Earnings are falling while revenue rises. The company is spending ahead of profit on the humanoid build-out, so the path to a sane multiple is not "grow into it" —

it is “grow into it while also dramatically expanding margin,” and there is no disclosed evidence the margin expansion is coming.

3. Samsung deployment is intent, not disclosed numbers. There are no public unit counts, no dated procurement schedule, no broken-out Samsung revenue. The single load-bearing element of the thesis is undated and unquantified. A 35% owner can also keep a captive supplier sub-scale indefinitely if that serves Samsung’s own P&L.

4. The technical setup is stretched into the print. RSI 71.6, +26.6% above the 50-day MA, with Q1 2026 earnings due May 15. The stock has run hard into a binary, the bar is high, and a print that shows no Samsung deployment revenue is a clean down-catalyst.

Optionality

Event	Date / window	Direction
Q1 2026 earnings	~May 15, 2026	Binary — first look for Samsung deployment revenue
First disclosed Samsung volume-procurement order	2026	Bull — converts intent to numbers
Bipedal humanoid product announcement with launch date	2026–2027	Bull if it comes; absence is a Gap
Samsung capital-discipline / memory-cycle pressure headlines	Ongoing	Bear — risk of captive R&D throttling
RB-Y1 collaborative variant deployed at scale in a Samsung fab	2026–2027	Bull on disclosure
Korea humanoid-theme sentiment rotation	Ongoing	Binary — drives the option premium both ways

---

The trade

Rainbow Robotics is a name where the thesis and the valuation point in opposite directions: the Samsung strategic anchor is genuinely rare and valuable, but the stock is already priced at 700x forward earnings as if that optionality has been cashed. This is a small, sentiment-driven position, not a conviction core holding. Initiate — if at all — at KRW 798,950–883,050 (current KRW 841,000 ± 5%), and be aware this is a melt-up-tape entry into a stretched technical setup (RSI 71.6, +26.6% above the 50-day MA) right before the Q1 2026 print on May 15. Size at 0.75% of risk capital — deliberately small; this is a Bucket C name carrying real chase risk and a valuation with nothing underneath it but Samsung's continued conviction. Stop at KRW 680,000, roughly 19% below current — a wide stop because the name is volatile and a tight one gets stopped on noise, but a hard stop because if the strategic-option premium starts to deflate it can deflate a long way. The catalyst is the Q1 2026 earnings around May 15, 2026 — the first read on whether any Samsung deployment revenue is appearing or whether the stock is still pure ownership-story. If you want exposure to the same "Korea wins humanoids" thesis with an owner that is committing harder and moving faster — building a 30,000-unit/year factory and committing its entire 2026 production to internal deployment — Hyundai Motor (005380) via the Boston Dynamics lens is the better-funded, faster-moving expression, and Hyundai is profitable. **Conviction: 5 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

SERV — Serve Robotics Inc. · WATCH (Tier-2) · Conv 5/10 · Bucket A

Serve Robotics Inc. (SERV)

Sidewalk delivery robots at real fleet scale and triple-digit revenue growth — burning roughly six dollars for every one it earns.

Investment Research · Photoncap-style deep dive · v1 of "Serve Robotics" · May 14, 2026

What Serve Robotics physically does

Serve Robotics operates a fleet of autonomous, four-wheeled, cooler-sized delivery robots that travel on sidewalks to carry restaurant and convenience orders the "last mile" — typically under a couple of miles — from a merchant to a consumer's door. The robot is a low-speed (walking-pace) electric vehicle with a lockable insulated cargo bay, a sensor stack of cameras, ultrasonics and lidar, and an onboard compute system running Level 4 autonomy within mapped service areas. It is purpose-built for one job: moving a small payload along a pedestrian path without a human driver and without a car.

The business model is the part that matters for the theme. Serve is not selling robots; it operates them and charges per delivery, increasingly blended with a software-services layer. In Q1 2026 software services contributed roughly one-third of revenue and just under half of total revenue was recurring, per the company's May 7, 2026 results release — a deliberate shift from "we own and run robots" toward "we also license the autonomy and operations stack." The binding-constraint argument for Serve is weak relative to the supply-chain names in this theme: sidewalk delivery is a service business with thin unit economics, not a chokepoint that everyone else depends on. Serve's edge, such as it is, is operational density and a multi-year head start on permitting, mapping and merchant integration in specific US metros.

The robot is general-purpose only within a narrow envelope — it does not manipulate, it does not enter buildings, it does not handle large or heavy loads. It is a logistics endpoint, and its economics live or die on deliveries-per-robot-per-day and the cost of the human remote-supervision overhead behind the "autonomous" fleet.

Product roadmap

Serve's hardware has progressed through generations of its sidewalk robot, with the current-generation unit (the third major design, scaled through 2025) being the platform underpinning the ~2,000-robot fleet referenced on the Q1 2026 call. The 2025 build-out was a fleet-expansion phase — the explicit Q1 2026 message was that the company has now "shifted from fleet expansion to increasing revenue per robot," per the May 7, 2026 release. That is the roadmap pivot: 2024-2025 was about putting robots on sidewalks; 2026 is about utilization.

The geographic roadmap: the operating footprint reached 44 cities across 14 states as of Q1 2026, driven by new metro launches, hospital-network deliveries (an indoor/campus extension of the model) and densification of existing markets. Cumulative deliveries across indoor and outdoor environments are approaching 2 million, per the company. The software roadmap — the more interesting one — is the productization of Serve's autonomy and fleet-operations stack as a licensable layer, which is what drove the software-services revenue to ~one-third of the Q1 mix.

What Serve does not do: it does not make humanoids, it does not do drone delivery, and it does not sell hardware to third parties at scale. It is a single-modality, single-use-case operator. The Uber relationship — Serve's roots trace to Uber, which remains a strategic shareholder and a demand channel via Uber Eats — is a distribution asset, not a product.

The financial print

Serve reported Q1 2026 on May 7, 2026: revenue of $3.0 million, up 238% sequentially and 578% year-on-year. The eye-catching line is that Q1 2026 revenue alone exceeded the company's entire FY2025 revenue — real growth, off a tiny base. Fleet revenue was nearly $2 million (roughly a 10x YoY increase) and software services made up the rest, with just under half of total revenue now recurring. The company reiterated FY2026 revenue guidance of $26 million and held non-GAAP operating expense guidance of $160-170 million for 2026.

That guidance pairing is the whole story: ~$26 million of revenue against ~$165 million of operating spend. Serve is spending roughly six dollars for every dollar it earns. Forward P/E is negative and not meaningful; this is a cash-burn-and-scale name. The offsetting fact is the balance sheet — Serve reported $197.4 million of liquidity at Q1 2026,

which at the guided burn rate funds roughly a year-plus of operations before another raise is needed. There is no named broad sell-side consensus of the Hanmi variety here; coverage is thin and the stock trades on fleet-growth and revenue-per-robot headlines.

One-year return: the stock has been volatile and is roughly flat-to-down over the trailing year, with RSI 47.8 and price 1.8% below the 50-day moving average — a neutral, unextended tape, unusual for this theme. Next earnings: Q2 2026, expected on or around August 6, 2026 — the binary on whether the "revenue per robot" pivot is actually lifting utilization.

Customer mix today

Serve's "customers" are best understood in two layers, and the company does not disclose a clean named-customer percentage breakout, so the figures below are reporting-basis. The demand layer is restaurants and merchants whose orders Serve fulfills — aggregated substantially through Uber Eats (Serve's largest demand channel and a relationship inherited from its Uber origins) plus a direct merchant network and, increasingly, hospital and campus clients. The revenue layer splits differently: roughly two-thirds fleet revenue (per-delivery economics) and one-third software services as of Q1 2026, per the May 7, 2026 release.

The structural shift worth flagging: in 2024-2025 Serve was almost entirely a per-delivery fleet operator; by Q1 2026 software services had grown to ~one-third of revenue and recurring revenue to just under half. That mix shift toward higher-margin, more predictable software is the most genuinely encouraging trend in the print — it is the difference between "low-margin gig logistics" and "autonomy platform." Whether it sustains is the open question; one quarter is not a trend.

What's actually happening at the merchant / Uber Eats demand layer

The mechanism that decides Serve's economics is deliveries-per-robot-per-day, and the Q1 2026 pivot language — "from fleet expansion to increasing revenue per robot" — is management conceding that utilization, not fleet size, is now the constraint. With ~2,000 robots deployed across 44 cities, the fleet is large enough; the question is whether each robot is busy enough to cover its capital and supervision cost.

The Uber Eats channel is the key demand pipe: it gives Serve order volume without

Serve having to build consumer demand, but it also caps Serve's pricing power and margin — Serve is a fulfillment subcontractor inside someone else's marketplace. The hospital-network and campus additions in Q1 2026 are an attempt to diversify into denser, more repetitive, higher-utilization environments where a robot can run more deliveries per shift. The honest read as of May 2026: Serve has proven it can deploy robots and grow revenue 578% YoY off a tiny base; it has not yet proven the per-robot economics work at scale, and the ~$165M-to-$26M spend-to-revenue ratio says the unit economics are nowhere near self-funding. The August 6, 2026 Q2 print is where "revenue per robot" either starts climbing visibly or it doesn't.

The competitive threat

Serve's competitive set is fragmented and the threat is less a single named rival than the structural question of whether sidewalk delivery is a defensible business at all. Direct sidewalk-robot competitors include Coco (backed by Sam Altman, operating in similar US metros), Starship Technologies (campus-focused, large cumulative delivery count), and Avride and others. None has obvious decisive scale over Serve, and Serve's ~2 million cumulative deliveries and 44-city footprint put it among the leaders by operational density.

The more serious threat is from above and from the side. From above: the delivery platforms themselves — Uber, DoorDash — could in-source or multi-source autonomy, and DoorDash has its own autonomous delivery efforts; Serve's dependence on Uber Eats as a demand channel is also a dependence on a party that could become a competitor. From the side: drone delivery (Zipline, Wing) attacks the same last-mile job with different economics, and conventional gig drivers remain the cheap incumbent. There is no material IP litigation driving the thesis. The competitive bottom line: Serve is a credible operator in a niche that may simply be a low-margin niche, and its biggest "competitor" is the possibility that no one makes good money doing this.

The terminal risk

The terminal risk for Serve is that autonomous sidewalk delivery is structurally a low-margin, capital-intensive service that never earns its cost of capital — not that a new technology obsoletes it, but that the business itself is the wrong shape. A sidewalk robot is a depreciating physical asset that needs maintenance, remote human supervision, permitting, mapping and insurance, to move a single small payload at walking pace. If

deliveries-per-robot-per-day plateaus below the level that covers all of that, scale makes the losses bigger, not smaller.

The secondary terminal risk is channel: if the delivery platforms in-source autonomy or squeeze fulfillment margins, Serve becomes a commoditized subcontractor. And the form-factor risk — drones for some routes, cheaper gig labor for others — means sidewalk robots may end up serving only a thin slice of last-mile geography. Serve has a credible roadmap toward the higher-margin software-services layer, which is the only real answer to this risk; if that layer scales, Serve becomes an autonomy platform rather than a robot operator. As of May 2026 that pivot is one quarter old. The multiple you can pay is constrained by the genuine possibility that the core service never reaches positive unit economics.

Bull / Gap / Optionality

Bull

1. Real, verified scale and growth. Q1 2026 revenue of $3.0 million was up 578% YoY and exceeded all of FY2025 revenue, with a ~2,000-robot fleet across 44 cities and nearly 2 million cumulative deliveries, per the May 7, 2026 release. This is no longer a science project — it is an operating fleet at meaningful scale, which is more than most early robotics names can claim.

2. The software-services mix shift is the right kind of growth. Software services reached ~one-third of revenue and recurring revenue just under half by Q1 2026. If Serve can keep shifting toward licensing its autonomy/operations stack, the margin profile improves structurally — this is the path from “gig logistics” to “platform.”

3. The balance sheet funds the experiment. $197.4 million of liquidity at Q1 2026 against ~$165 million of guided 2026 opex gives roughly a year-plus of runway. Serve is not at a financing cliff, and the un-extended tape (RSI 47.8, slightly below the 50-day MA) means you are not chasing.

4. Distribution via Uber Eats is a genuine asset. Serve gets order volume without building consumer demand, and Uber remains a strategic shareholder. For an early operator, having a demand pipe already plumbed is a real de-risking of the go-to-market.

Gap

1. The spend-to-revenue ratio is brutal. FY2026 guidance pairs ~$26 million of revenue with ~$160-170 million of opex — roughly six dollars spent per dollar earned.

Even with triple-digit growth, the path to self-funding is years away and depends on utilization economics that are unproven at scale.

2. Utilization, not fleet size, is now the binding constraint — and management said so. The Q1 2026 pivot language ("from fleet expansion to increasing revenue per robot") is an admission that the ~2,000-robot fleet is under-utilized. Until deliveries-per-robot-per-day visibly climbs, the fleet is a cost, not an engine.

3. Channel dependence cuts both ways. Reliance on Uber Eats for demand caps Serve's pricing power and exposes it to a partner that could in-source autonomy or become a competitor. Serve is a subcontractor inside someone else's marketplace.

4. The niche may simply be low-margin forever. Sidewalk delivery faces drones from one side and cheap gig labor from the other, and no operator in the category has demonstrated attractive unit economics. The terminal risk is that the business is the wrong shape, not that Serve executes it badly.

Optionality

Event	Date / window	Direction
Q2 2026 earnings	~August 6, 2026	Binary — does revenue-per-robot pivot show up?
FY2026 revenue guide ($26M) confirm/raise	Through 2026	Bull if raised; Bear if cut
Software-services mix trend	Each quarterly print	Bull if >1/3 and rising
New large channel / merchant partnership	2026	Bull — diversifies off Uber Eats dependence
Next equity raise	~2027 (runway-dependent)	Bear — dilution overhang

---

The trade

Serve is a small, speculative, watch-it-prove-the-economics name, and the trade reflects that. Entry zone is current ±5%, roughly $8.51–$9.41 — and the tape genuinely cooperates here, with RSI 47.8 and price slightly below the 50-day moving average, so this is one of the few names in the theme you are not chasing into a melt-up. Size small, around 0.75% of risk capital, treating SERV as a high-variance call option on sidewalk-delivery

economics rather than a core holding; the spend-to-revenue ratio means a wrong thesis here is a slow bleed, not a sudden gap. Stop at roughly $7.10, below recent structural support. The defining catalyst is the Q2 2026 print around August 6, 2026 — specifically whether the “revenue per robot” pivot translates into visibly rising utilization, because fleet scale is already proven and utilization is the only thing left to prove. If the thesis is “autonomous last-mile logistics is real,” the cleaner expression is a diversified or platform-layer exposure rather than a single-modality sidewalk operator; Serve is the purest sidewalk-robot bet available, with all the concentration risk that implies. **Conviction: 5 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

454910 — Doosan Robotics (454910.KS) · SKIP / WAIT (Tier-3) · Conv 4/10 · Bucket C

Doosan Robotics (454910.KS)

Asia’s leading cobot brand pivoting to an AI-systems-integrator and humanoid story — strong franchise, shrinking revenue, roadmap still on paper.

Investment Research · Photoncap-style deep dive · v1 of “Doosan Robotics” · May 14, 2026

What Doosan Robotics physically does

Doosan Robotics makes collaborative robot arms — “cobots” — and is one of the largest cobot brands in Asia. A cobot is a lightweight, force-limited robot arm designed to work safely alongside humans without a safety cage: it senses contact and stops, so it can be deployed on a factory floor next to a worker rather than fenced off the way a traditional industrial robot arm is. Doosan’s cobot range spans payloads from a few kilograms up to the heavier 20kg-plus class, where the company has historically been particularly strong — by some 2022 estimates it held roughly 72% of the global cobot market for payloads of 20kg or more. The use cases are machine-tending, palletising, assembly, welding, screwdriving, food service (Doosan’s coffee-and-food robot cells are a visible product) and inspection — the work that sits between fully manual and fully automated.

The technology proposition is deployment speed and flexibility. A traditional six-axis industrial robot from Fanuc or Yaskawa is faster and more precise, but it requires integration engineering, safety fencing and fixed installation. A Doosan cobot is meant to be set up quickly, redeployed across tasks, and programmed by a non-specialist. That has made cobots the fast-growing segment of factory automation and the natural on-ramp for small and mid-sized manufacturers who cannot justify a full automation cell.

The 2025–2026 strategic story is Doosan trying to move beyond being a cobot hardware vendor. The company is repositioning as an AI-powered “intelligent robot solutions” provider and systems integrator — selling modular, fast-to-deploy automation solutions rather than just arms — and has publicly committed to humanoid robot development. It opened the Doosan Robotics Innovation Center in Bundang in September 2025 to

consolidate its AI, software and humanoid R&D, and stood up dedicated AI/software and humanoid development teams in the second half of 2025. Doosan does not yet make a shipping humanoid; the humanoid line is a stated strategic intent backed by an R&D centre, not a product. Doosan also does not make its own AI foundation models — like its Korean peers, the intelligence layer is the build-or-partner challenge ahead.

Product roadmap

The shipping product line is the cobot range — the H-, M-, A- and E-series collaborative arms across the payload spectrum, plus packaged solutions like the food-and-beverage robot cells. Doosan has cumulatively shipped over 10,000 cobot units as of late 2025, a real installed base. The product cadence here is incremental: payload extensions, controller and software upgrades, easier programming, and the "intelligent robot solutions" packaging that bundles the arm with vision, software and integration into a faster-deploy product.

The roadmap event that re-rated the stock is the humanoid and AI pivot. The confirmed milestones: a strategic shift announced in 2025 toward "practical humanoids" capable of complex tasks; the formation of dedicated AI/software and humanoid R&D teams in H2 2025; the opening of the Doosan Robotics Innovation Center in Bundang in September 2025; and a "Best of Innovation" award at CES 2026 for the company's intelligent-robot-solutions positioning. What is *not* on the roadmap is a humanoid product with a name, a launch date and a price. As of May 2026, Doosan's humanoid effort is an R&D programme and a strategic narrative, not a product line — the distinction the company's valuation does not currently make. Treat the cobot installed base and shipment numbers as confirmed; treat the humanoid roadmap as stated intent only, with no dated deliverable.

The financial print

This is where Doosan Robotics gets uncomfortable. FY2025 revenue was approximately KRW 32.98 billion (~$23M USD) — *down* roughly 29.6% from KRW 46.83 billion in 2024. The cobot business shrank materially in 2025. The company has been loss-making — forward P/E is not meaningful because the company is not expected to print a profit on the near-term consensus horizon, which is why the frontmatter forward_pe is recorded as 0.0 (not applicable). Trailing-twelve-month revenue as of late 2025 was around $22M USD, consistent with the shrinking top line. Against that, the market capitalisation is

roughly KRW 6.9 trillion (~$4.9B USD) — a company doing ~$23M of *declining* revenue, losing money, valued near $5 billion.

The honest framing: Doosan Robotics is valued almost entirely on the humanoid/AI-integrator narrative and on its strategic position within the Doosan Group, not on its printed financials. The earlier attempt to merge Doosan Robotics into the much larger, profitable Doosan Bobcat — at a stock-swap ratio that was widely criticised for over-valuing the loss-making Robotics unit — was cancelled in December 2024 after Korean market turmoil and shareholder opposition; that episode is itself evidence that even the controlling group struggled to find a defensible valuation for the standalone Robotics business. There is no robust multi-broker FY2026 earnings consensus to anchor to, because the earnings are negative and the story is narrative-driven; treat any forward figure as estimate-basis.

The binary event is the Q1 2026 earnings, due around May 15, 2026 — the print that shows whether the FY2025 revenue decline has bottomed and reversed, or whether the cobot core is still contracting while the humanoid story remains pre-product. With the stock +16.2% above its 50-day MA and RSI at 61.5 going in, there is some momentum cushion but also a real “show me” bar.

Customer mix today

Doosan Robotics’ customer mix is its genuine strength and is the opposite of the customer-concentration risk that afflicts a name like UBTECH or Rainbow. Doosan sells cobots through a broad global distribution network — distributors and system integrators reselling to small and mid-sized manufacturing end-users across North America, Europe and Asia. The company opened a US headquarters to push North American distribution and has built one of the more extensive cobot channels in the industry. There is no single dominant customer; the revenue is diversified across thousands of end-users and a wide distributor base. Doosan does not disclose a top-customer concentration percentage because the structure is genuinely diffuse — that diffuseness is real, and it is a positive.

But the structural shift in the mix is the *wrong* direction on volume. FY2025 revenue fell ~29.6%, meaning the diversified channel sold materially fewer cobots than in 2024 — a function of soft global manufacturing capex, intensifying price competition from lower-cost cobot makers, and a destocking cycle in the distributor channel. The humanoid “customer mix” does not exist yet — there are no humanoid customers because there is no humanoid product. So the customer picture is: a healthily diversified but shrinking cobot base, and a humanoid customer base of zero. The valuation prices a future

humanoid customer mix that has not begun to form.

What’s actually happening in the cobot channel

The mechanism that matters for the near term is whether Doosan’s cobot revenue decline is cyclical or structural. The cyclical read: 2025’s global manufacturing capex was soft, SMEs deferred automation spend, and distributors worked down inventory — a destocking cycle that mean-reverts when manufacturing capex recovers, and the Q1 2026 print is where a recovery would first show. The structural read: cobots are commoditising. Lower-cost Chinese cobot makers are competing hard on price, the technology is maturing to the point where the hardware is increasingly undifferentiated, and a Doosan cobot now competes not only with Universal Robots and Fanuc’s CRX line but with a widening field of cheaper alternatives. A 29.6% revenue decline in a single year is too steep to be purely cyclical — some of it is structural pricing and share pressure.

This is precisely why Doosan is pivoting to “intelligent robot solutions” and AI-integration: the company is trying to escape commoditised hardware competition by moving up the stack into software, vision-integrated solutions and faster modular deployment, where margin and differentiation are defensible. The CES 2026 “Best of Innovation” award and the Bundang Innovation Center are the visible artefacts of that pivot. The honest assessment in May 2026: the pivot is strategically correct but unproven in the numbers — there is no disclosed evidence yet that the solutions/AI mix is offsetting the hardware decline. The Q1 2026 print is the first real test of whether the channel has bottomed and whether the up-the-stack move is generating revenue.

The competitive threat / Universal Robots and the Chinese cobot field

Doosan’s direct competitive set on its actual revenue base: Universal Robots (the Danish company, the global cobot market leader with the most comprehensive portfolio and global presence — the benchmark Doosan is measured against), Fanuc’s CRX collaborative line (Fanuc bringing its industrial-robot scale, reliability and 8-year-maintenance-free positioning into the cobot segment), Stäubli, and a broadening field of lower-cost Chinese cobot makers competing aggressively on price. Industry evaluations place Universal Robots as the market leader and Doosan as an “emerging leader” — strong, expanding, but not the category leader. On the humanoid side, Doosan competes — on paper — with Hyundai’s Boston Dynamics, Rainbow Robotics (its Samsung-backed Ko-

rean peer), UBTECH, Unitree and Tesla. On humanoids Doosan is the furthest behind of any of these: it has an R&D centre and a press release where the others have products.

The competitive bear case is the squeeze: in cobots Doosan is being pressured from below by cheap Chinese hardware and from above by Universal Robots' scale and Fanuc's reliability brand, and its revenue is shrinking as a result. In humanoids it is behind everyone with a product. The competitive bull case is narrower — Doosan has a real, large, trusted cobot brand and installed base, a genuine global channel, and Doosan Group backing, and if the AI-solutions pivot works it could leverage that channel to sell higher-value integrated systems rather than commoditised arms. There is no active IP litigation of note involving Doosan Robotics as of May 2026. The competitive risk is not legal — it is being structurally squeezed in the business it actually has, while being structurally behind in the business its valuation prices.

The terminal risk

The terminal risk for Doosan Robotics is being caught between two losing positions. On one side, cobots get commoditised: the hardware becomes a low-margin, Chinese-cost-structure-dominated category, and Doosan's cobot revenue and margin erode structurally rather than cyclically — the 29.6% FY2025 decline is the leading edge of that, not a one-off. On the other side, the humanoid pivot never produces a competitive shipping product: Doosan starts the humanoid race years behind Boston Dynamics, UBTECH and Unitree, without Samsung-scale backing of the kind Rainbow has, and the "practical humanoid" stays an R&D programme. If both happen, Doosan is a shrinking cobot maker with a humanoid story that never ships, and the KRW 6.9 trillion valuation has no support.

The transition timing that governs the downside is shorter than for the pure humanoid names because the cobot-commoditisation pressure is happening *now*, in the printed revenue, not in some 2030s scenario. The humanoid-parity timing is the same long-tail question as for the rest of the theme. The named alternative beneficiaries of Doosan's squeeze are Universal Robots and Fanuc (taking cobot share with scale and reliability) and the components supply chain (selling into whoever wins). Doosan does have a credible asset — the brand, the 10,000-unit installed base, the global channel — and the AI-solutions pivot is the strategically right response. But "strategically right" and "valued correctly" are different things, and at ~$5B for a shrinking, loss-making, ~$23M-revenue business, the multiple is pricing a turnaround and a humanoid franchise that are both still hypothetical.

Bull / Gap / Optionality

1. Genuine cobot franchise — brand, installed base, global channel. Doosan is one of Asia's leading cobot brands with over 10,000 units shipped and one of the industry's more extensive distribution networks, including a US headquarters built out for North American reach. This is a real asset that a humanoid startup cannot replicate — a trusted brand and a route to market.

2. The AI-solutions pivot is the strategically correct response to commoditisation. Moving up the stack from commoditised hardware to AI-integrated "intelligent robot solutions" and modular systems integration is exactly the right move — it is where margin and differentiation live. The CES 2026 "Best of Innovation" award and the September 2025 Bundang Innovation Center are tangible commitment, not just slides.

3. Cyclical revenue decline could mean-revert. Part of the FY2025 29.6% revenue drop is a soft-manufacturing-capex and distributor-destocking cycle. If global manufacturing capex recovers through 2026, the cobot channel restocks and the top line can bounce — the Q1 2026 print is the first place a recovery would show.

4. Doosan Group backing and strategic intent. Doosan Robotics sits inside the Doosan Group, which has the balance sheet to fund a multi-year AI and humanoid R&D push. The group's strategic commitment — dedicated humanoid teams, a consolidated innovation centre — means the humanoid effort is funded, not starved.

5. Cobots are still a structurally growing category long-term. Even with near-term price pressure, collaborative robots remain the fastest-growing segment of factory automation as SMEs automate. If Doosan defends share and rides the category, the installed base compounds — the FY2025 decline notwithstanding.

Gap

1. Revenue is shrinking — down ~29.6% in FY2025. This is the single most important fact and the valuation ignores it. A ~$23M-revenue company doing *less* revenue than the year before, valued near $5 billion, is being priced on narrative entirely. A decline this steep is too large to be purely cyclical.

2. The humanoid roadmap is a press release, not a product. As of May 2026 there is no Doosan humanoid with a name, a launch date or a price — only an R&D centre and strategic intent. Doosan is the furthest behind of any humanoid-narrative name in this batch, and it lacks the Samsung-scale captive backing that gives Rainbow its story.

3. Squeezed from both sides in the core business. Cheap Chinese cobot makers compete from below; Universal Robots' scale and Fanuc's reliability brand compete

from above. Doosan is an “emerging leader,” not the category leader, and its shrinking revenue is the evidence of the squeeze.

4. Even the controlling group couldn’t agree a valuation. The cancelled 2024 Doosan Bobcat merger — abandoned partly because the swap ratio was widely seen as overvaluing the loss-making Robotics unit — is direct evidence that pricing this business is contentious even for insiders. The standalone valuation rests on a turnaround that has not started in the numbers.

Optionality

Event	Date / window	Direction
Q1 2026 earnings	~May 15, 2026	Binary — has the cobot revenue decline bottomed?
First named humanoid product with launch date	2026–2027	Bull if it ships; absence is a deepening Gap
Evidence AI-solutions mix is offsetting hardware decline	2026	Bull on disclosure
Global manufacturing capex recovery data	2026	Bull — drives cobot channel restock
Doosan Group strategic action (new merger/restructuring)	Ongoing	Binary — group-level moves swing the valuation
Chinese cobot pricing / share-take data	Ongoing	Bear — structural commoditisation evidence

---

The trade

Doosan Robotics is the weakest risk/reward in this batch and the position reflects it — a small, tactical holding at most, not a conviction name. The franchise is real, but the revenue is shrinking, the humanoid roadmap is still a press release, and the ~$5B valuation prices a turnaround and a humanoid business that are both hypothetical. If you take a position, initiate at KRW 101,460–112,140 (current KRW 106,800 ± 5%) — a melt-up-tape entry, with the stock +16.2% above its 50-day MA and RSI at a not-yet-extended 61.5. Size at 0.5% of risk capital — the smallest sizing in the batch; this is a Bucket C name where the chase risk is high and the fundamental support is thin. Stop at

KRW 88,000, roughly 18% below current, beneath the structural support — a wide stop because the name is volatile and narrative-driven, but a hard one because if the Q1 print confirms the cobot decline is structural rather than cyclical, the narrative premium can unwind sharply. The catalyst is the Q1 2026 earnings around May 15, 2026 — the test of whether the FY2025 revenue collapse has bottomed. The cleaner expression of the same thesis: if you want Korean cobot/robotics exposure, Rainbow Robotics (277810) at least has a real strategic owner committing capital and a shipping semi-humanoid product, and if you want the broader robotics theme, the industrial-arm incumbents Fanuc (6954) and Yaskawa (6506) are profitable, growing, and gaining cobot share rather than losing it. **Conviction: 4 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

RR — Richtech Robotics Inc. · SKIP / WAIT (Tier-3) · Conv 3/10 · Bucket B

Richtech Robotics Inc. (RR)

A $623M micro-cap built on ~$5M of service-robot revenue, a $329M cash pile, and a humanoid named “Dex” that has not shipped.

Investment Research · Photoncap-style deep dive · v1 of “Richtech Robotics” · May 14, 2026

What Richtech Robotics physically does

Richtech Robotics is a US-listed service-robotics company that builds and deploys functional robots for hospitality, food service and similar commercial environments — robot baristas and bartenders, delivery and bussing robots for restaurants and casinos, and related autonomous service hardware. The robots are wheeled, fixed-task machines: a robotic arm that makes a drink, a tray-carrying mobile base that runs food across a venue floor. They are useful, demonstrable, real products — but they are not a choke-point technology and not a binding constraint on anything. This is a service-robot integrator and operator, and that framing is important for the theme: Richtech sits at the commoditized, low-barrier end of robotics, not the high-margin supply-chain end.

The company is mid-transition in business model. Historically Richtech sold robots as hardware; it is now pushing a Robotics-as-a-Service (RaaS) model — leasing the robots and charging recurring fees, plus an emerging “data services” layer. In Q1 FY2026 (the quarter ended December 31, 2025) RaaS revenue was just $0.3 million, up 31% YoY — tiny in absolute terms but the line management wants investors to watch. Gross margin on the business is high (reported around 76%), which tells you the hardware is being recognized at healthy markups, but the absolute revenue base is so small that margin percentage is close to noise.

The newer, narrative-driving piece is “Dex,” a humanoid robot Richtech showcased at CES 2026, described as powered by NVIDIA’s Jetson Thor compute module and aimed at dynamic commercial environments. Dex is the reason RR appears in a humanoid screen at all. As of May 2026 Dex is a showcase unit, not a shipping product.

Product roadmap

Richtech's established product lines are its service robots: the ADAM robotic barista/bartender system, the Scorpion beverage robot, the Medbot and Titan-class delivery/service robots for hospitality and casino floors, and related autonomous service units deployed across US venues. These are the revenue base today — hardware sales plus the nascent RaaS leases.

The roadmap event that matters is Dex. Richtech unveiled Dex at CES 2026 (January 2026) as a humanoid built on NVIDIA Jetson Thor, and on its Q1 FY2026 commentary (results dated mid-February 2026) the company said the Dex humanoid rollout is "expected in 2Q FY2026" — i.e. the quarter ending roughly June 2026. Treat that as a target, not a confirmed ship date; Richtech has a pattern of CES showcases that lead reveals by some months. The strategic framing management offers is a transition "toward recurring RaaS and data services revenue streams," with Dex as the higher-capability platform that extends the RaaS model from fixed-task service robots into general-purpose manipulation.

What Richtech does not have: scale, a proven humanoid, a named anchor customer of consequence, or any component/supply-chain position. It is a downstream operator competing in two crowded categories at once — commodity service robots and frontier humanoids — with the resources of a micro-cap.

The financial print

Richtech reported Q1 FY2026 (quarter ended December 31, 2025) in mid-February 2026: total revenue of roughly $5.0 million, up about 19% YoY, with RaaS revenue of $0.3 million (+31% YoY) and a GAAP net loss of $8.4 million — but that loss was driven by $8.3 million of non-cash stock-based compensation, so the cash operating loss is far smaller than the headline. Gross margin was reported around 76%. The standout balance-sheet item is total liquidity of $328.8 million — an enormous cash pile relative to a company doing ~$5 million of quarterly revenue and carrying a ~$623 million market cap.

That cash-to-revenue mismatch is the defining financial fact. Richtech has raised far more capital than its business currently consumes or its revenue justifies, which means two things: there is no near-term financing cliff (a genuine positive for a micro-cap), but the share count has expanded heavily to build that pile, and the $8.3 million of quarterly

stock-based comp signals ongoing dilution as a structural feature, not a one-off. Free cash flow margin has been reported deeply negative.

There is no meaningful named sell-side consensus for FY2026 — coverage is minimal and the stock trades on news flow (CES showcases, Dex updates, customer announcements) rather than on a modeled forward number. Forward P/E is negative and not informative. The 1-year tape: the stock is volatile, currently with RSI 58.8 and price 18.7% above the 50-day moving average — extended on the short term, consistent with momentum chasing the Dex narrative. Next earnings: the next quarterly print is expected around late May 2026, roughly May 31, 2026 — the binary on whether Dex actually rolled out on schedule.

Customer mix today

Richtech does not disclose a clean named-customer percentage breakout, so this section is necessarily reporting-basis and thin. The deployed base is hospitality, restaurant, food-service and casino operators using Richtech's service robots — robot baristas in venues, delivery/bussing robots on restaurant and gaming floors. The company periodically announces individual customer or venue wins, but no single customer is disclosed as a dominant share of revenue, and the absolute revenue (~$5 million a quarter) is small enough that any one mid-size deployment moves the mix.

The structural shift management is trying to engineer is from hardware-sale customers to RaaS-lease customers — recurring relationships rather than one-time box sales. As of Q1 FY2026 that shift is barely underway: RaaS revenue is $0.3 million of a ~$5 million total, so roughly 94% of revenue is still non-recurring. The honest read is that Richtech's "customer mix" is too small and too undisclosed to analyze with the rigor this theme's larger names allow — and that thinness is itself the finding.

What's actually happening at the end market

Because there is no single dominant customer, the relevant question is what is happening in Richtech's end markets — hospitality and food-service automation — and whether the RaaS pivot plus Dex can convert that into a durable business. The hospitality service-robot market is real but brutally competitive and price-sensitive: venue operators adopt a robot barista or a delivery robot if the payback is fast and the maintenance burden is low, and they switch vendors easily because the products are substitutable. Richtech's 19% YoY revenue growth in Q1 FY2026 shows demand exists; the 76% gross margin

shows it can sell at a markup; the $0.3 million RaaS line shows the recurring-revenue transition is real but embryonic.

The Dex bet is an attempt to escape the commodity service-robot trap by moving up to general-purpose manipulation. But Dex enters a humanoid race where Tesla, Figure AI, Boston Dynamics, Agility and a wave of well-funded Chinese players are spending vastly more. Richtech's advantage, if it has one, is that it already operates robots in real commercial venues and has customer relationships and an operations stack — so it could deploy a humanoid into known environments rather than starting cold. That is a thin edge. The proof point is the 2Q FY2026 Dex rollout: either it ships into real venues with disclosed economics, or it remains a CES showcase. As of May 2026 it is the latter.

The competitive threat

Richtech faces competition on two fronts simultaneously, which is the core problem. In service robots, the field is crowded with established and low-cost players — Bear Robotics (backed by SoftBank, large restaurant-robot deployments), Pudu Robotics and Keenon Robotics (high-volume Chinese service-robot makers with aggressive pricing), and numerous regional integrators. These competitors are larger, ship more units, and compete on price in a category with low switching costs. Richtech's ~$5 million quarterly revenue is small against them.

On the humanoid front, Dex competes — at least notionally — against Tesla's Optimus, Figure AI's Figure 03, Boston Dynamics' Atlas, Agility's Digit and China's Unitree and UBTech. Every one of those programs is better capitalized for the actual R&D, even though Richtech's $329M cash pile is large relative to its own size. There is no material IP litigation driving the thesis. The competitive bottom line: Richtech is sub-scale in its core market and a long-shot entrant in its aspirational one. The bull case is not that Richtech wins either race outright — it is that the cash pile and the existing venue relationships give it more optionality than a $623M micro-cap usually has.

The terminal risk

The terminal risk for Richtech is twofold and structural. First, the service-robot business is a commoditized hardware category — robot baristas and delivery bots are increasingly substitutable, made cheaply at volume by Chinese competitors, and sold into price-sensitive venue operators with low switching costs. A commodity hardware business with a sub-scale player is a value trap; the RaaS pivot is the intended escape, but

at $0.3 million of recurring revenue it is unproven. Second, Dex enters the humanoid race years behind and orders of magnitude under-resourced relative to the leaders; the realistic outcome for Dex is not “Richtech wins humanoids” but “Dex is a niche service-humanoid that may or may not find an economic foothold.”

There is no clean technology-transition framing here because Richtech is not the incumbent in any category that something newer would obsolete — it is the small challenger in two categories at once. The honest terminal risk is mundane: the company never reaches the scale where its economics work, the cash pile is slowly consumed by stock-comp-heavy operations, and the equity dilutes its way down. The multiple you can pay is constrained by the fact that ~$5 million of quarterly revenue does not support a $623 million market cap on fundamentals — the valuation is a narrative-and-cash-pile valuation, not an earnings valuation.

Bull / Gap / Optionality

Bull

1. The cash pile removes the financing cliff. $328.8 million of liquidity at Q1 FY2026 against a business that burns a fraction of that in cash terms means Richtech can fund years of operations and the Dex program without an imminent raise. For a micro-cap, runway is the rarest asset, and Richtech has an unusual amount of it.

2. Real revenue, real margins, real growth. Q1 FY2026 revenue of ~$5.0 million was up ~19% YoY at a ~76% gross margin. Richtech is not pre-revenue — it has shipping products, paying customers and a healthy markup, which puts it ahead of pure-concept robotics names.

3. The RaaS transition is the right strategic direction. RaaS revenue up 31% YoY, even from a $0.3 million base, plus an emerging data-services layer, is the correct pivot away from commodity hardware sales toward recurring revenue. If it scales, the business becomes structurally more valuable.

4. Existing venue relationships give Dex a warm-start deployment path. Unlike a pure humanoid startup, Richtech already operates robots in real commercial environments. If Dex ships, it can be placed into known venues with known operations support — a thin but genuine edge over cold-start competitors.

Gap

1. The valuation is unsupported by fundamentals. A $623 million market cap on ~$5 million of quarterly revenue is a ~30x-plus revenue multiple on a sub-scale operator in commoditized categories. The valuation rests on the Dex narrative and the cash pile, not on earnings power.

2. Heavy, structural dilution. The $8.3 million of stock-based compensation in a single quarter — larger than the headline net loss — signals that share-count expansion is an ongoing feature of how Richtech operates and funds itself. The big cash pile was bought with equity, and more will be.

3. Sub-scale in a price-competitive core market. Bear Robotics, Pudu and Keenon ship far more service robots and compete on price. Richtech's ~$5 million quarterly revenue is small, and the service-robot category has low switching costs and thin defensibility.

4. Dex is a showcase, not a product. As of May 2026 Dex has been demonstrated at CES but not shipped. The "2Q FY2026 rollout" is a target from a company with a pattern of showcase-ahead-of-reveal timing, entering a humanoid race led by players spending vastly more.

Optionality

Event	Date / window	Direction
Dex humanoid rollout	"2Q FY2026" (~quarter ending June 2026)	Binary — ships into venues or stays a showcase
Next quarterly earnings	~May 31, 2026	Binary on revenue trend + Dex timing
RaaS revenue trajectory	Each quarterly print	Bull if RaaS grows toward a meaningful mix
Named anchor customer announcement	2026	Bull — would validate the deployment model
Capital deployment of the $329M pile	2026-2027	Binary — accretive use vs. slow burn

---

The trade

Richtech is a speculative micro-cap and the trade is sized as a lottery ticket, not a position. Entry zone is current ±5%, roughly $2.65–$2.93, but with the explicit caveat that the tape is already extended — RSI 58.8 and price 18.7% above the 50-day moving average mean you would be buying into a momentum move on the Dex narrative, not at a base. Size very small, around 0.4% of risk capital or less, treating RR as a high-variance call option on the Dex rollout and the RaaS pivot; the fundamentals — ~$5 million quarterly revenue against a $623 million cap — do not support more conviction than that. Stop at roughly $2.20, below recent structural support, accepting that a micro-cap of this kind can gap through a stop. The defining catalyst is the cluster around late May 2026 — the next quarterly print and the 2Q FY2026 Dex rollout window — which together test whether Dex is a product or a press release. If the thesis is "humanoids and service robotics are real investable themes," the cleaner expression is the better-capitalized OEMs and the supply-chain names in this theme; Richtech is the most speculative, lowest-conviction way to get the exposure. **Conviction: 3 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

KITT — Nauticus Robotics, Inc. · SKIP / WAIT (Tier-3) · Conv 2/10 · Bucket E

Nauticus Robotics, Inc. (KITT)

A distressed subsea-autonomous-robot micro-cap on a Nasdaq compliance monitor, kept breathing by a financing line — a survival story, not yet an investment.

Investment Research · Photoncap-style deep dive · v1 of "Nauticus Robotics" · May 14, 2026

What Nauticus Robotics physically does

Nauticus Robotics builds autonomous subsea robots — underwater vehicles designed to do offshore work without the umbilical tether and dedicated support vessel that conventional remotely operated vehicles (ROVs) require. The flagship platform is Aquanaut, a free-swimming subsea robot that can transit autonomously and then carry out intervention tasks — inspecting, manipulating and servicing underwater infrastructure — using onboard autonomy and manipulator arms rather than a human pilot driving it through a cable from a ship on the surface. The supporting software stack, branded ToolKITT (the source of the ticker), is the autonomy and operations layer.

The thesis the company sells is genuinely sensible in the abstract: the offshore energy and subsea-infrastructure industry spends enormous sums on vessel time, because conventional ROVs need a large, expensive ship parked overhead for the duration of a job. A robot that can be deployed from a smaller, cheaper vessel — or eventually a subsea docking station — and work autonomously attacks that cost structure. Offshore wind, oil and gas, telecom cables and, increasingly, defense and survey work all need subsea inspection and intervention. Nauticus is targeting that "robotics replaces vessel-hours" cost arbitrage.

But for this theme the framing has to be brutally honest. Nauticus is not a binding-constraint supplier and not a scaled operator — it is a distressed micro-cap with an $8.6 million market cap, fighting Nasdaq compliance issues, that has executed a reverse stock split and is being kept alive by an external financing commitment. The product

idea is reasonable; the company's survival is the actual question, and everything else is secondary to it.

Product roadmap

Nauticus's product roadmap centers on Aquanaut — the autonomous subsea vehicle — moving from demonstration and qualification toward commercial offshore deployment. The 2026 cadence of disclosures is essentially a series of readiness milestones: in April 2026 Nauticus announced it had completed Simulated Intervention Testing to advance "East Coast offshore readiness," and it disclosed an award for an offshore archaeological investigation project along the US East Coast. The company has also announced an Aquanaut manufacturing-expansion intention and "accelerated growth strategy" language, plus a planned UAE expansion tied to a strategic-investor relationship. The ToolKITT software layer is the autonomy roadmap — the differentiation Nauticus claims over tethered ROVs.

The honest read on the roadmap: these are small, project-level and readiness-level milestones, not a product-generation roadmap with confirmed volume commitments and dates. There is no disclosed multi-unit fleet rollout schedule, no named anchor customer committing to a deployment program, and the "manufacturing expansion" is an intention from a company that does not currently have the balance sheet to fund a large build. Distinguish clearly: the East Coast survey award and the simulated intervention test are confirmed events; the manufacturing expansion, the accelerated growth and the UAE expansion are stated strategy, not contracted backlog.

What Nauticus does not have: scale, a fleet generating recurring revenue, a confirmed large customer program, or the financial capacity to execute an ambitious roadmap without further external capital.

The financial print

The financial print is where the thesis confronts reality, and the reality is distress. Nauticus disclosed 2025 results showing higher revenue but a large net loss — revenue is minimal in absolute terms and nowhere near covering the cost structure. The defining financial facts as of spring 2026 are structural, not operational: Nauticus executed a 1-for-8 reverse stock split effective April 21, 2026 (a classic move to keep a share price above Nasdaq's minimum bid requirement), and on April 27, 2026 received notice from

Nasdaq confirming compliance with continued listing requirements — but with conditions. The company remains subject to a Mandatory Panel Monitor and must maintain a minimum shareholders' equity of $3.5 million for each fiscal quarter through December 19, 2026. That is a company on probation with its own exchange.

The lifeline is an external financing commitment: a strategic investment agreement for up to $50 million with Master Investment Group, tied to a UAE expansion. That facility — to the extent it funds — is what stands between Nauticus and a going-concern crisis. There is no meaningful forward consensus, no named sell-side coverage of substance, and forward P/E is not applicable (no earnings, and the metric is meaningless for a company at this stage). The market cap is approximately $8.6 million — a number that itself tells you the equity market is pricing this close to a wipeout scenario.

The price action confirms the distress: the reference price is $1.98, RSI 28.5 (deeply oversold), and price 49.3% below the 50-day moving average — a near-total collapse relative to recent levels. One flag worth noting explicitly: with a 1-for-8 reverse split effective April 21, 2026 and a stock down ~49% versus its 50-day MA, the moving-average and RSI figures span the split and should be read as "the stock has collapsed," not as a precise technical signal. Next quarterly results are expected around mid-2026 — and the binary there is less "what was revenue" and more "did the company stay above the $3.5 million equity floor and is the financing facility funding."

Customer mix today

Nauticus does not have a customer mix in any meaningful institutional sense, and that thinness is the finding. Revenue is project-based and minimal: the most concretely disclosed recent customer event is the award of an offshore archaeological investigation project along the US East Coast. There is no disclosed dominant recurring customer, no multi-year service contract of consequence, and no percentage breakout to construct — the revenue base is too small and too episodic.

The end markets Nauticus targets — offshore energy (oil and gas, offshore wind), subsea telecom, survey work, and potentially defense — are real and large. But targeting a market and having customers in it are different things, and Nauticus is closer to the former. The "customer mix" section for a healthy name in this theme breaks out named customers by percentage; for Nauticus the honest version is that the company is still at the stage of winning individual projects to prove the platform works offshore, not the stage of having a customer base to analyze. The structural "shift" that would matter — the first repeatable, multi-unit customer program — has not happened.

What’s actually happening — the survival fight, not a customer mechanism

For a healthy name this section dissects the share-gain mechanism at the dominant customer. For Nauticus there is no dominant customer and no share-gain mechanism; what is “actually happening” is a fight to remain a going concern and stay listed, and that is what an investor must underwrite.

The concrete state of play as of May 2026: Nauticus is on a Nasdaq Mandatory Panel Monitor through December 19, 2026, must hold shareholders’ equity above $3.5 million every fiscal quarter, has just done a 1-for-8 reverse split, and is relying on a financing commitment of up to $50 million from Master Investment Group to fund operations and a planned UAE expansion. The company has been actively reshaping its leadership — a Chief Revenue Officer appointment effective May 13, 2026, a new General Counsel, and other “leadership enhancements” — which can be read two ways: a genuine commercialization push, or the kind of management churn typical of a company under existential pressure. The operational milestones (the April 2026 simulated intervention test, the East Coast survey award) are real and modestly encouraging, but they are small relative to the financial hole.

The honest assessment: the question for KITT is not “will it gain subsea market share” — it is “will the financing fund, will it stay above the equity floor, and will it stay listed.” Every product and customer consideration is downstream of that. As of May 2026 the answer is genuinely uncertain, which is why the equity is priced at an ~$8.6 million market cap.

The competitive threat

In the abstract subsea-autonomy market, Nauticus faces a credible competitive set: established subsea-vehicle and ROV providers such as Oceaneering International (the dominant offshore ROV services player, vastly larger and financially stable), Saab’s Seaeye subsea division, Kongsberg’s underwater-robotics business, and a field of autonomous-underwater-vehicle specialists. Several of these are pushing their own untethered and autonomous capabilities — the “robot replaces vessel-hours” idea is not proprietary to Nauticus, and the large incumbents have the balance sheets and the customer relationships Nauticus lacks.

But naming competitors slightly overstates where Nauticus is in its lifecycle. The real competitive threat to Nauticus is not that Oceaneering takes its share — it is that Nauticus runs out of money and time before it ever competes at scale. There is no material IP

litigation driving the thesis. The competitive bottom line: in a market where the credible players are large, well-capitalized incumbents with established offshore customer relationships, a sub-$10-million-market-cap company on an exchange compliance monitor is not a competitive participant in any practical sense — it is an aspirant trying to survive long enough to become one.

The terminal risk

For most names in this theme the terminal risk is a technology transition that obsoletes a product family. For Nauticus the terminal risk is far more immediate and far more mundane: insolvency and delisting. The company must hold shareholders' equity above $3.5 million every quarter through December 19, 2026 under a Nasdaq Mandatory Panel Monitor; it has just reverse-split to maintain its bid price; and it depends on an external financing facility funding as promised. A missed equity-floor test, a financing facility that does not fully fund, or a continued cash burn that the facility cannot cover all lead to the same place — a delisting and a probable equity wipeout. That is the terminal risk, and it dominates everything else.

The product-level structural risk — that large incumbents like Oceaneering and Kongsberg productize autonomous subsea intervention faster and better, leaving Nauticus's Aquanaut without a defensible niche — is real but almost academic given the financial situation. A company has to survive to be obsoleted. The multiple an investor can pay is not constrained by a technology transition; it is constrained by the fact that the equity may be worth zero within twelve months. This is not a name where you weigh a thesis against a multiple — it is a name where you size for total loss.

Bull / Gap / Optionality

Bull

1. The underlying idea is genuinely sound. Autonomous, untethered subsea robots that reduce expensive vessel-hours attack a real cost problem in offshore energy, wind, telecom and survey. If Nauticus survives and the platform proves out, the addressable market is large and the value proposition is logical — the concept is not the problem.

2. There is a financing lifeline. The strategic investment agreement for up to $50 million with Master Investment Group, tied to a UAE expansion, is a real external capital

commitment — the difference between an immediate going-concern crisis and a fighting chance. To the extent it funds, it buys time.

3. Nasdaq compliance was confirmed — for now. The April 27, 2026 notice confirmed compliance with continued listing requirements (albeit conditionally, under a monitor). The company is not delisted; it is on probation, which is a step better than the alternative.

4. Concrete operational milestones are landing. The April 2026 simulated intervention test and the East Coast archaeological survey award are real, confirmed events showing the platform doing actual offshore-relevant work — modest evidence the technology is more than a slide deck.

Gap

1. This is a distressed company fighting for survival. A Nasdaq Mandatory Panel Monitor through December 19, 2026, a mandatory $3.5 million quarterly equity floor, a 1-for-8 reverse split in April 2026, and an ~$8.6 million market cap. The equity market is pricing this close to a wipeout, and the financial facts justify that pricing.

2. The financing lifeline is a commitment, not cash in the bank. "Up to $50 million" is a ceiling on a facility, contingent on terms and conditions. If it does not fully fund, or funds with heavy dilution, the survival case weakens immediately — and dilution at this market cap is severe for existing holders.

3. There is no real customer base or recurring revenue. Revenue is minimal and project-episodic; there is no disclosed anchor customer program. The company is at the prove-the-platform stage, not the scale stage, while burning cash it does not comfortably have.

4. The competitive field is full of large, stable incumbents. Oceaneering, Kongsberg, Saab Seaeye and others have the balance sheets, customer relationships and offshore track record Nauticus lacks — and they are pursuing autonomous subsea capability too. Even in the survival case, Nauticus competes from far behind.

Optionality

Event	Date / window	Direction
Next quarterly results / $3.5M equity-floor test	~Mid-2026	Binary — survival-critical
Master Investment Group facility funding	Through 2026	Binary — the lifeline either funds or it doesn't

Event	Date / window	Direction
Nasdaq Panel Monitor period end	December 19, 2026	Binary — clears the monitor or faces delisting risk
New offshore project awards	2026	Bull — modest evidence of commercial traction
Further dilution / capital raises	2026	Bear — severe at this market cap

---

The trade

Nauticus is not an investment in the institutional sense — it is a lottery ticket, and it must be sized and described as one. If a position is taken at all, the entry zone is current ±5%, roughly $1.88–$2.08, but the entry-zone framing is almost a formality: this is a binary survival bet, not a name where 5% on the entry matters. Size at 0.25% of risk capital or less — a true lottery-ticket allocation where the base case is total loss and the position exists only to capture the small probability that the financing funds, the company clears its Nasdaq monitor, and the subsea-autonomy idea gets a chance to prove out. Stop at roughly $1.40, while accepting that a distressed micro-cap can gap straight through any stop on a delisting headline or a financing failure. The defining catalyst is the next quarterly results and the associated $3.5 million shareholders'-equity compliance test around mid-2026 — the binary is not earnings quality, it is whether the company stays alive and listed. If the thesis is "autonomous subsea robotics is a real long-term market," the overwhelmingly cleaner expression is a financially stable incumbent — Oceaneering or Kongsberg — that will still be standing in three years; Nauticus is the distressed, maximum-risk, probably-zero way to express it, and should be treated accordingly. **Conviction: 2 / 10.**

Sources referenced inline throughout. Reference v1 of this template format: `_Watchlist/hanmi-photoncap-style.md`.

Section 8 — How the layers interconnect

The five layers are not independent. A robot is a control loop: information flows up the stack (senses → brain), commands flow down (brain → body), and materials feed in from the side (magnets → motors). In a single second of a humanoid's operation, the loop runs thousands of times. Understanding the connections is essential because it explains why a breakthrough at one layer cascades to demand at the others.

The data-flow loop, worked example: humanoid picks up a mug. The vision SoC (Ambarella CVflow or Lattice FPGA glue layer) processes camera feeds, identifies the mug and computes its 3D pose; the lidar (Ouster ToF or Aeva FMCW) confirms the distance. The brain (NVIDIA Jetson Thor running GR00T) decides the grasp approach — come in from the right, fingers spaced 80 mm apart, target the handle. The motion planner (also on Jetson, lower-rate) breaks the plan into joint trajectories. The joint controller (Ambiq MCU at 1 kHz) translates trajectory into motor currents, ms by ms. Each motor receives current; the magnet field rotates; the strain-wave reducer (Harmonic Drive in the wrist) steps RPM down 100:1 and torque up 100×; the joint moves; bearings (Minebea Mitsumi) absorb the load. The encoder (Allegro magnetic chip) reports joint position back to the controller. The force sensor (VPG strain-gauge) reports the contact force when fingers touch the mug. If the force is too high, the controller eases off in the next millisecond. The loop closes. New camera frame. New plan. Repeat ~30 times per second for the visual loop, 1,000+ times per second for the joint control loop.

Every layer is on the critical path. A breakdown in any one — a chip that runs too hot, an encoder that drifts, a reducer with too much backlash, a magnet that demagnetises — and the robot fails. This is why "good enough" is not a position in any of these layers: the joint layer has to be Japanese-precision good; the brain has to be real-time good; the senses have to be photo-realistic and low-latency good; the magnets have to be NdFeB-strong. Substitutes do not exist at most layers.

The BOM-cost cascade. A roughly $35-50k humanoid in 2026 distributes BOM cost approximately as follows: joints (Layer 1) ~30%; battery + power electronics (not in this

universe) ~15%; sensing and perception (Layer 2) ~12%; compute and memory (Layer 3) ~10%; structural frame, harnesses, assembly ~25%; magnets and other materials (Layer 4) ~8%. The 2026 BOM is roughly half the 2024 BOM, driven by joint-layer cost reductions (China cost-down on reducers, motor commoditisation in the lower-tier markets) and compute commoditisation at the cortex tier (Jetson Thor is the standard). The implication is that **the dollar-volume growth in the supply chain is being carried by unit growth, not by per-unit ASP**. Unit volumes need to ~2× per year for the next three years for the supply-chain absolute revenue thesis to underwrite the multiples in this basket. Helix-02 and the Schaeffler-Humanoid AI commitment are the two most concrete signals that this is happening.

Inter-layer correlated trades. Three cross-layer pair-trade structures fall naturally out of the cascade:

- *Magnet-to-motor pair:* MP Materials (Layer 4) and Nidec or Minebea (Layer 1) trade as a connected pair. A NdFeB price spike helps MP but hurts the motor makers' input costs; a NdFeB price collapse does the opposite. The MP / 6479 pair captures the layer arbitrage.
- *Brain-to-glue pair:* NVIDIA (Layer 3 cortex) and Lattice (Layer 3 visual cortex) are explicit collaborators (NVIDIA Halos AI Systems lab with Lattice FPGAs). The two names co-move on humanoid platform-design wins. The pair captures the platform-vs-glue value distribution.
- *Sense-to-platform pair:* Mobileye (Layer 2 perception) and Tesla (Layer 5 OEM) are not direct trading pairs — Mobileye is not in any Tesla BOM — but the broader perception-platform-versus-vertical-OEM trade is structurally cleanest expressed as MBLY versus a humanoid pure-play like 9880 UBTECH. The pair captures the "perception-as-a-platform" thesis at the layer boundary.

Section 9 — The 38-name shortlist by layer

The full universe at a glance. Each row carries the same data the dashboard reads from `_Dashboards/coverage.json`. Tier within each layer is sorted by conviction descending. Prices are May 14 2026 yfinance pulls.

#	Ticker	Name	Layer
Layer 4 — Materials (1 name)			
1	MP	MP Materials	Materials
Layer 1 — Body / joints (8 names)			
2	6324	Harmonic Drive Systems	Body
3	6268	Nabtesco	Body
4	6479	MinebeaMitsumi	Body
5	SHA	Schaeffler AG	Body
6	6481	THK	Body
7	NOVT	Novanta	Body
8	ALNT	Allient	Body
9	6594	Nidec	Body
Layer 2 — Senses (10 names)			
10	CGNX	Cognex	Senses
11	ALGM	Allegro MicroSystems	Senses
12	6861	Keyence	Senses
13	OUST	Ouster	Senses
14	MBLY	Mobileye	Senses
15	AEVA	Aeva Technologies	Senses
16	ST	Sensata Technologies	Senses
17	INDI	indie Semiconductor	Senses
18	VPG	Vishay Precision Group	Senses
19	ARBE	Arbe Robotics	Senses
Layer 3 — Brain (5 compute + 1 memory)			

#	Ticker	Name	Layer
20	NVDA	NVIDIA	Brain
21	QCOM	Qualcomm	Brain
22	LSCC	Lattice Semiconductor	Brain
23	AMBA	Ambarella	Brain
24	AMBQ	Ambiq Micro	Brain
25	MRAM	Everspin Technologies	Memory
Layer 5 — OEMs (13 names)			
26	SYM	Symbotic	OEM (specialt
27	005380	Hyundai Motor	OEM (strategi
28	6954	Fanuc	OEM (industri
29	6506	Yaskawa Electric	OEM (industri
30	TSLA	Tesla	OEM (strategi
31	TER	Teradyne	OEM (specialt
32	ABB	ABB Ltd	OEM (industri
33	9880	UBTECH Robotics	OEM (humanc
34	277810	Rainbow Robotics	OEM (humanc
35	SERV	Serve Robotics	OEM (specialt
36	454910	Doosan Robotics	OEM (humanc
37	RR	Richtech Robotics	OEM (specialt
38	KITT	Nauticus Robotics	OEM (specialt

 = upgrade vs May 14 master, on PhotonCap May 19 read.

Section 10 — Action sheet (May 22 2026)

The actionable trades, in priority order. Buckets reference `references/price-discipline.md`: A = entry-zone-actionable now; B = modestly above; C = significantly above (chase risk); D = missed; E = better than reference.

★★★ Top single action: MP Materials (MP)

The only Tier-1 BUY in this basket currently at Bucket A — RSI 44, ~flat to its 50-DMA, undistorted by the broader robotics-momentum tape. Entry zone $56.76–$62.74 (current $59.75 within zone). Stop $51. Sizing 2.0% of risk capital (held position is already ~25K shares on the JPM book). The thesis is structural: NdFeB rare-earth magnet bottleneck for every humanoid robot; only US integrated mine-to-magnet producer; Apple/GM/DoD backstops; 2026 metal-and-magnet plant ramp. PhotonCap mention count 13× corroborates.

★★ Buyable today (Bucket A or close)

- **Symbotic (SYM) — $48.56, Bucket A, RSI 34.** Oversold. GAAP-profitable warehouse robotics integrator with $22B+ backlog. Risk is Walmart concentration (~84%+), not the tape. Entry $46-51. Stop $40. Sizing 1.5%.
- **Schaeffler (SHA0.DE) — €9.91, Bucket B, RSI 73 ↑ upgrade.** 14× forward earnings, Q2 2026 first series production of the humanoid actuator platform, four named OEM customers, Humanoid AI 5-year supply contract. Entry €9.41-€10.41. Stop €8.30. Sizing 1.5%.
- **Hyundai Motor (005380.KS) — W712,000, Bucket B, RSI 79.** 14× automaker that contains Boston Dynamics. Pullback target W620-680k for a cleaner entry; current price still actionable for a smaller scale-in. Sizing 1.0% scale-in, 2.0% full at pullback.

- **Allegro MicroSystems (ALGM) — $45.44, Bucket B, RSI 55.** Best risk/reward in sensing. Magnetic position-sensing leader, two named humanoid joint design-ins, robotics revenue line doubled YoY, cheapest growth multiple in the layer, calmest chart. Entry $43-47. Stop $36. Sizing 1.5%.
- **Lattice Semiconductor (LSCC) — $123.63, Bucket B, RSI 58.** Strongest pure-robotics compute thesis: low-power FPGA sensor-fusion glue layer, NVIDIA Halos and TI partnerships, AI mix already in the print, relatively calm tape. Entry $115-130. Stop $100. Sizing 1.5%.

★ Pullback queue (Bucket C/D — wait for entry)

- **Harmonic Drive (6324.T), Nabtesco (6268.T), MinebeaMitsumi (6479.T)** — the Japanese joint-layer oligopoly. All Tier-1 BUY by thesis, all extended (RSI 69-84). Pullback targets: 6324 to ¥5,600; 6268 to ¥5,200; 6479 to ¥3,500.
- **VPG (Vishay Precision) — RSI 90, +95% vs 50-DMA.** Thesis is right (humanoid tactile sensing leader) but entry is non-investable here. Wait for a 25-35% pullback to $70-78.
- **Tesla (TSLA) — Bucket C.** Optimus option value inside a 176× stock. Own for the company, not the robot.
- **AMBQ, MRAM — Bucket D.** Both extraordinarily extended (+116% / +174% vs 50-DMA). Deep-pullback only: AMBQ $54-62, MRAM $29-34.

✗ Skip / avoid

- **Nidec (6594.T)** — accounting scandal; uninvestable until clean restated accounts. Tier-3 SKIP, revisit only after restated FY25 results.
- **Doosan Robotics (454910)** — FY2025 revenue declined ~30%; humanoid roadmap a press release; ~$5B cap on declining revenue.
- **KITT, RR, ARBE** — speculative micro-caps; lottery tickets if any sizing at all.
- **ZBRA Zebra Technologies** — out (exited robotics by selling Fetch Robotics to Skild AI).

Section 11 — Risk overlay

The 38-name basket carries five structurally distinct risks. None is dominant in isolation; the interaction is what matters.

(1) Tape extension. The universe is broadly extended — many names print RSI 70-90 as of May 14 2026. Per the workspace two-layer framework, RSI extension inside a secular theme is a timing tool, not an exit trigger. But VPG (RSI 90, +95%), AMBQ (RSI 90, +116%) and MRAM (RSI 78, +174%) are in a separate valuation-and-momentum category; entry discipline on those is non-negotiable. A 10-20% momentum unwind across the basket is the most plausible near-term risk.

(2) Humanoid timeline slip. Optimus has missed multiple internal targets between 2023 and 2025. Figure AI's Helix-02 livestream is the strongest reliability proof to date, but commercial-volume deployment in 2026-2027 is still gated. A high-profile programme delay (Tesla downsize, Figure capital crunch, BYD pull-back) triggers a sector de-rate.

(3) Chinese cost competition compressing the joint layer. Leaderdrive ("Green Harmonic"), Estun and Inovance are the named A-share challengers to Harmonic Drive and Nabtesco. Our exclusion of these names on the IBKR constraint means our universe is *long the incumbent* and *short the disruptor by absence*. The pricing-power thesis at the joint layer rests on the incumbent share defence holding through 2027.

(4) Rare-earth supply politics. China processes ~85% of NdFeB. Export controls in 2024-2025 already disrupted global prices. The MP Materials thesis is partly a hedge against this; but if Chinese export-control normalisation occurs (unlikely but not impossible), MP loses its scarcity premium.

(5) OEM margin disappointment. Humanoid OEMs are mostly pre-profit or loss-making. The supply chain is monetising on order books that may or may not convert to full deployment volumes. If an OEM cohort posts unit-economics that imply secular gross-margin disappointment, the order-book confidence at the supplier layer collapses faster than the underlying physical demand.

The five risks are partially correlated. A momentum unwind (1) is most likely to be

triggered by humanoid timeline slip (2) or OEM margin disappointment (5). Chinese cost competition (3) is the slowest-moving but most structurally important risk for the joint-layer Tier-1 BUYs. Rare-earth politics (4) is the most binary and the hardest to handicap.

Section 12 — Companion deliverables and how to use them

This document sits in a constellation of robotics deliverables produced over May 14-22, 2026. The reader should know what each one is for:

Deliverable	Path	What it is	When to use it
`Robotics-Universe-by-Layer-v1-DEEP-DIVE.pdf`	workspace root	**This document.** Layer-organised navigation map + 38 per-name deep dives + PhotonCap May 19 integration	Starting point for any new robotics decision; "where does this name sit in the stack?"
`Robotics-MASTER.md` / `Robotics-DEEP-DIVE.pdf`	workspace root	Tier-organised version (BUY / WATCH / SKIP) of the same 38 per-name dives	"What can I act on right now?" — sorted by conviction
`_Watchlist/tradingview-import-robotics-v1.txt`	`_Watchlist/`	TradingView import — paste once into "Add multiple symbols" to create the watchlist sections	Setting up the chart workspace
`_Watchlist/tradingview-color-cheatsheet-robotics-v1.md`	`_Watchlist/`	Colour-flag cheat sheet (GREEN BUY now / YELLOW WATCH / RED SKIP) with exchange-picker resolution table	After pasting the TV import — set the colour flags

Deliverable	Path	What it is	When to use it
`Robotics-Deep-Dive-v2.pptx` / `.pdf`	workspace root	47-slide executive deck with 8-slide Section 0 primer	Showing the thesis to anyone unfamiliar with the stack
`Theme -- Robotics/` folder	workspace root	Per-ticker subfolders for ongoing news/earnings/Bloomberg accumulation	Ongoing maintenance — drop new docs into the relevant ticker folder
`_Dashboards/coverage.json`	`_Dashboards/`	The 27-field contract — one entry per ticker	What every dashboard reads from

The maintenance protocol is the same as for the AI Universe v2 document: when a new dated catalyst lands (next earnings, partnership announcement, IPO filing), update the per-name file in `Theme -- Robotics/<TICKER>/<ticker>-deep-dive.md` and re-run `update_dashboard.py` to refresh the coverage entry. When mention counts in the weekly Substack digest cross a threshold (Schaeffler 46× on PhotonCap May 19 is the canonical example of this trigger firing), re-tier the affected names and re-render the master.

The Robotics dashboard build, which this document is the prerequisite for, will read `_Dashboards/coverage.json` filtered on `theme == "Robotics"` and present the same data in interactive form. The dashboard is not built yet — that is the next step after sign-off on this document.