

MACRO OPS

TRADE · LEARN · EVOLVE

THE LONG PULL



Podcast Notes: Astrid Wilde

By Brandon Beylo

September 09, 2026



This week we have a special *Podcast Notes* series from my conversation with Astrid Wilde. Astrid is the founder of Inheritance HQ. According to its website, the company “transforms observation into training data that robots and world models can learn from at scale, enabling capability to compound instead of restart with every new embodiment.”

Don't worry if you didn't understand a word from that sentence. I didn't either. Which is why I wanted to interview Astrid.

We're deeply curious (and bullish) about Physical AI and the Robotics thematic at MO over the next 18-24 months. However, I only had “outside” knowledge of the theme. In other words, I digested a bunch of bullish Big Investment Bank reports, but still craved the “boots on the ground” primary contact to pressure-test these assumptions.

Astrid is that guy.

Besides founding a Physical AI/robotics data company, Astrid is a great public markets investor. We spend a lot of time exploring his investment process, why doing less has equaled more (read: higher returns), and how Astrid cultivated a strong network of people to feed him ideas.

I hope you guys enjoy this piece. I learned a ton from Astrid, and this *Podcast Notes* report, powered by Claude, does a great job highlighting everything we discussed.

Let's get after it.

The Human Prior: A Crash Course in Robot Data

Wilde's framing of the physical AI landscape starts with a distinction most public-market narratives skip. In his telling, two fundamentally different design architectures share the label “robot.”

The first is deterministic — a dishwasher, a factory arm bolted to an assembly line, a Roomba on a schedule. It does one thing, under known conditions, with a fixed set of rules. The second class increasingly captures capital and attention: machines expected to handle edge cases in unstructured environments.

"That's the perfect robot. But there are so many cases where the task is more complicated than just doing one thing with a deterministic set of rules. And we need more quote-unquote intelligence baked in." — Astrid on the split between deterministic and intelligent robotics

For the second class, Wilde invokes what has become AI's guiding principle: the "bitter lesson." Complex problems are no longer solved by writing better rules. They are solved by providing a model with enough varied, high-quality data and letting the attention mechanism do the work.

Inheritance AI exists because that lesson, once you accept it, converts the humanoid robotics race into a data-supply problem — and data supply, unlike compute, is not something the hyperscalers can simply write a check for.

Primitives Beat End-to-End Tasks

The most economically important insight in the interview may be Wilde's point about primitives. It is tempting, when thinking about training a robot to (say) play tennis, to imagine capturing full matches end-to-end.

Wilde argues that is exactly the wrong instinct. A tennis rally contains dozens of things a player is doing simultaneously — tracking score, positioning relative to an opponent, reading spin, choosing strategy. But the primitive motor behaviors that actually need to be captured are much smaller: standing on two feet, tracking a ball, swinging forehand, swinging backhand, adjusting for spin at contact.

Capturing the primitives is dramatically cheaper than capturing the full task, and because the same primitives recur across dozens of downstream problems, it generalizes better. This is the economic logic of Inheritance AI's data pipeline.

While competitors pay people to wear cameras for eight-hour shifts, Wilde's team thinks harder about the atomic actions a well-tuned model can recombine on its own.

"There are so many tasks where it's economically inefficient for us to capture the actual full task end-to-end, but where it's very easy for us to capture the primitives." — Wilde on why primitives dominate his data strategy

The Messy Reality of Human Data

The romantic version of the "camera on the head" data-collection approach glosses over the grunt work that makes it usable. Wilde is blunt about what an eight-hour unedited video of a factory worker actually contains: bathroom breaks, five-minute chats with coworkers, and — his phrasing — "people pulling up something obscene and watching it while they work."

None of that is training data. All of it has to be labeled, segmented, and filtered before a model ever sees it.

The pipeline that turns raw human capture into "meaningful segments" is, in Wilde's framing, the real bottleneck and the real moat.

Anyone can stick a GoPro on a warehouse worker. Very few teams have the tooling and the labor to convert that footage into the crisp, scored, primitive-level examples that a large model can actually learn from.

The supply-chain squeeze no one is talking about

The most actionable piece of physical-AI intelligence in the interview may be Wilde's account of what has happened to **sensor supply over the past eighteen months**.

Eighteen months ago, he explains, a small robotics team could email a sensor or camera vendor and get three people on a half-hour call, a single test unit, and delivery within a week. Today, the same vendors will not even take the meeting **unless the order is measured in hundreds or thousands of units**.

"Every form of sensor is starting to sell out. We're getting estimates of, oh, we'll have more in three months, oh, we'll have more next year. That's not a problem we had eighteen months ago." — Wilde on the tightening physical AI supply chain

This has not yet worked its way into public-market narratives the way GPU or memory tightness has. It is a leading indicator worth watching, because it feeds directly into the next wave of humanoid, robotics, and industrial-automation product launches.

If sensor lead times keep extending, the timing of every physical-AI product roadmap slides with them — and the pricing power of the very few upstream sensor manufacturers goes up.

How training environments actually work

One of the most technically rich exchanges came when Wilde walked through how a team actually trains a physical AI model to perform a novel task — say, controlling a robot arm to draw an image.

The answer, as he tells it, is deceptively simple: build an environment in Isaac Sim or MuJoCo, express the task in a form the model can score itself against, and let the model iterate.

The clever piece is the encoding. A drawing task feels like something a language model has no business doing — until you remember that vector image formats are, quite literally, text. A model can be trained to "draw" by emitting SVG-style vector text.

Alternately, a vision model can be trained to score pixel-by-pixel against a target image. Either way, an abstract creative problem gets reduced to a deterministic, scoreable computation, and that is enough for the model to learn on if you give it enough environments and enough compute.

"There's no magic sauce or secrets under the hood. It's all very very simple stuff. It's just so many of them, and so much compute thrown at the problem, that eventually all of the connective tissue kind of works." — Wilde on the mechanics of reinforcement-learning environments

The unresolved question, Wilde is quick to add, is "how many is enough." The industry has been consistently wrong — in his own experience, wrong by orders of magnitude — about how much environment data and reinforcement-learning signal is required to unlock a new capability. That demand side of the story keeps hyperscalers building.

Inheritance AI's scope, and where it does not compete

Wilde is careful about where Inheritance AI plays and where it does not. His models are tuned to recognize human action, so anything that starts with "here is video of a human doing X" is in-bounds.

Forklifts and other human-augmented devices — where the operator is not physically performing the manipulation, but driving a machine that is — are explicitly out of scope. That crispness matters for anyone trying to size the addressable market or map

Inheritance AI's competitive set: this company is built around human proprioception and manipulation, not vehicle telemetry.

The niche-robotics future

I pushed on the humanoid-versus-niche debate — is physical AI really a two-horse race between Tesla's Optimus and Figure, or is the interesting surface area much more fragmented?

Wilde's answer was one of the more optimistic notes in the conversation: **the near future belongs to a proliferation of specialized machines**, not to a small number of generalized humanoids.

He points to Matic — the modern robot vacuum that quietly handles edge cases the original Roomba never could — as a template.

The winning form factors will be almost invisible: silent, purpose-built, doing one job extraordinarily well across a variety of conditions. That is the shape of the "return time to people" thesis he keeps coming back to. Menial labor gets automated one small chip at a time, and the compounding effect over five to ten years is enormous.

"I think that you will be very pleasantly surprised about the very near future then, Brandon." — Wilde on the case for specialized, niche robotics over humanoids

Public Markets: Doing Less to Capture More

Wilde's public-market process, in his own framing, has exactly two components.

The first is what he calls "front-running earnings"—not in the quant sense of forecasting a beat by ten cents, but in the broader sense of identifying a material, durable change in a business or sector before consensus catches up. The second is patience: sit around and wait for the first thing to happen.

"The thing that I'm looking for is what are things that are going to materially change in a business before consensus has caught up to it. And thing two is just sit around and wait for thing one to happen." — Wilde on his two-part public-market process

The historical wins

For the last three to four years, this framework has led Wilde into what he now shorthands as "the AI trade" — but he was there earlier and in more idiosyncratic form than the label suggests.

He bought crypto miners before they pivoted to high-powered compute. He bought power assets before the market recognized power as the binding constraint on data-center growth. Last year he bought memory names before the industry admitted it was sold out for four years.

Each of these was, at the time of purchase, a variant view about a sector that most investors had either written off or not connected to the AI capex cycle.

The insight that reshaped his process

The single most useful observation may be Wilde's admission about the marginal value of his own work.

From roughly 2019 through 2021, he was, in his phrasing, "a private equities analyst" — reading earnings releases, consuming news, spending six to ten hours a day on the market. Today he spends "maybe fifteen minutes a day, optimistically" on public markets. His compounded return has not changed.

"The number of hours that I spend reading earnings releases and the news has no material impact on my performance. So sincerely — I mean this without a smidgen of insincerity — find a new hobby." — Wilde on effort and outcomes in public equities

This is what I said in response: Investing is a rare domain where increased effort doesn't translate into better outcomes. In tennis, in software engineering, in almost every craft, six hours a day of practice for two years produces obvious improvement. In public markets, the same six hours a day produces about the same result as fifteen minutes, and often worse, because it invites the constant reshuffling of positions that erodes long-term returns.

Wilde's edge isn't that he works harder than the average investor; **it is that he has structured his life to notice rare, large opportunities and do nothing the rest of the time.**

The portfolio, decoded

As of September 4, 2026, Wilde's publicly disclosed positions above 1% included Nvidia, QXO, Nintendo, SK Hynix, Amazon, and a couple of undisclosed positions, alongside a meaningful cash sleeve. Read against his stated philosophy, each holding is doing specific work:

- **Nvidia and SK Hynix** are the direct expressions of the ongoing "front-run the AI capex cycle" thesis — the memory position in particular ties back to Wilde's point that leading hyperscalers require a 20%-plus internal return on data-center builds, which will keep both the compute and memory supply chains structurally tight.
- **Nintendo** is doing something more subtle. Wilde does not frame it as a variant-view growth story; he explicitly calls it a "cash position that also hedges my memory exposure." Nintendo, unlike most hardware companies of its scale, has historically not locked in long-term memory contracts and instead buys on the spot market. That leaves it unusually sensitive to memory prices — the mirror image of an SK Hynix position — which makes it a natural offset.
- **Cash** is not a residual for Wilde; it is a position. It reflects the "wait for thing one" side of the process. He does not force allocation into ideas that do not clear his bar.

Nebius: the case study

The single cleanest illustration of Wilde's process is his account of the Nebius trade. The setup was almost accidental: a portion of a Russian-linked company listed in the U.S., and its shares simply appeared in the brokerage accounts of anyone who had previously held the Russian ETF or related exposure. Forced holders reflexively clicked "sell," regardless of what the underlying business actually was.

Underneath the forced selling was a real, sector-relevant cloud-compute business — one Wilde already had context on through his AI-adjacent network. That collision — a theme he understood, a business he could verify, and a counterparty he could name — produced what he described as one of the most attractive risk/reward setups he expects to see in his lifetime.

"When you're in a situation in public markets where you know who your counterparty is, you're probably going to make a lot of money. And if you have a materially different

view — in the case of data centers being this long-standing trend — you can make phenomenal life-changing money." — Wilde on the Nebius trade and identifying the other side

Wilde's checklist is basically:

- (1) Is it something I understand?
- (2) Has it been pre-screened by people I trust who follow the same themes I do?
- (3) Do I know the counterparty and their motivation?
- (4) Do I have a materially different view from the market?

That's it.

Where he is currently looking

Wilde flagged two open themes he has not yet fully sized. The first is Builders First Choice (BLDR) — a direct bet on the U.S. interest-rate path and on home construction, expressed today through a starter position in QXO. The second is specific consumer companies the market still prices as damaged goods from the COVID hangover.

Peloton (PTON) and GameStop (GME) are the two names he raised, though he doesn't hold either.

On Peloton, the stock is down more than 96% from its COVID high; if the business has stabilized at all, it may be trading at a single-digit forward earnings multiple. On GameStop, Wilde says he "can't believe I'm saying this" but that a rough annualization of last quarter's figures implies an EV/EBIT under three.

He is not calling either a buy — he explicitly declined to take a position — but he is flagging both as names where the historical hype has scared serious investors away and left a variant-view opportunity for anyone willing to underwrite the underlying business quality.

The GameStop framing is especially interesting. Wilde is not underwriting a return to the video-game retail model. He is asking whether GameStop is quietly becoming a physical hub for the collectibles market — Pokémon cards, sports cards, the broader graded-collectibles boom — and whether management's cost base and store footprint could support that pivot. Whether or not that is the right read, it is the right shape of question.

Why public-market robotics is a trap right now

The most notable omission in Wilde's portfolio is direct physical-AI exposure. That is deliberate, and his reasoning matters for anyone building a robotics book today. The public options, he argues, are almost all derivative bets: FANUC and ABB make real robots, but everything else is a bet that "the demand from physical AI use cases is going to materially lead to inflection in their business," and you are simultaneously taking on the operational and management risk of each specific company.

"The further away you get from the bet you're interested in making, the worse your risk-return prospects kind of look. All of the most interesting stuff happening in robots is still in private markets." — Wilde on the limits of public-market physical AI exposure

For an investor with a strong robotics view, Wilde's advice is essentially to either go direct through private-market access or to wait for an IPO cycle that gives clean public exposure. Neither is comfortable, but both are more honest than pretending a derivative bet is the same trade.

Commodities: A Regime Change & Ingenuity Discount

Wilde is explicit that resource demand — particularly in the U.S. and Europe — has undergone a **material regime change relative to the steady-to-declining trend that held from roughly the 1980s through the mid-2010s**. That reversal, in his view, still isn't priced into many second- and third-derivative names: oilfield service companies, coal shippers, LNG infrastructure, and industrial suppliers to the broader energy complex.

"The trends we've seen since the 1980s of resource demand — specifically in the U.S. and Europe — being kind of steady to down is completely reversed. And I don't think that's a material regime change that continues to be priced into all number of things." — Wilde on the underappreciated resource-demand regime change

Why he avoids direct commodity exposure

Despite that view, Wilde also wants to avoid direct, long-term commodity exposure. His reasoning is worth sitting with. Any sustained price increase, he argues, invites profit-seeking behavior — capital floods the sector, capacity comes on, and innovation quietly lowers the marginal cost of extraction. The Permian is his running example: the

industry has repeatedly unlocked resources and refining processes that prior generations of engineers thought were uneconomic or physically infeasible.

The pattern is not unique to oil. It shows up in every metals cycle Wilde has watched: high prices attract capital, capital funds innovation, innovation compresses the cost curve, and the cycle rolls over.

The way to play the regime change, in his view, is through companies servicing extraction — where operational leverage compounds — rather than through direct commodity exposure that human ingenuity will eventually arbitrage away.

Automation and mining: a five-to-ten-year setup

I asked Astrid whether underground mining tasks — drilling, blasting, hauling — could become genuinely automatable, and whether the resulting cost reductions could offset demand growth and cap commodity prices.

Wilde's answer was direct: yes, and probably within a five-to-ten-year window. Prices first have to rise enough to motivate serious capital into the automation problem; once motivated, human ingenuity, in his phrasing, will solve it.

"Never ever find yourself betting against human ingenuity." — Wilde on the durability of commodity price spikes

The interesting angle here is that this automation wave is still primarily a private-market story. Wilde noted that a meaningful cohort of startups has been funded over the last two years explicitly targeting the mining stack — resource discovery, extraction, and transportation.

Public-market equivalents (Sandvik, Caterpillar, and similar) are, again, derivative bets. This is likely the next frontier where private allocators will earn structurally better exposure than public-market participants until an IPO cycle unlocks direct comparables.

The commodity-bull's counter

The consensus view among natural-resource investors is that the easy material has already been mined — that grades are declining, that discovery pipelines are thin, and that all the technology in the world will not compensate for the geologic reality.

Wilde's counter is not that the geology has changed, but that the technology stack applied to it has advanced enough to make previously uneconomic material economic. Both statements can be true simultaneously; the investable question is which effect dominates on any given commodity, on any given time frame.

For copper and uranium — the two commodities most exposed to the AI and electrification capex cycle — the demand pull is currently overwhelming the ingenuity discount. That may not be true forever, and Wilde's framework is a useful reminder of the mean-reverting force that eventually shows up.

AI Investment, Demand, & National Security

The most intellectually honest moment in the interview came when I asked what he had failed to raise. Wilde's answer was the question he cannot resolve for himself: what is the marginal demand for higher intelligence?

If today's open-source models are already "good enough" for the vast majority of tasks — and Wilde suspects they are — then much of the future capex on training and inference may be malinvestment.

Every additional order of magnitude of compute produces diminishing marginal utility for the median consumer, who is already receiving something functionally close to general intelligence. Wilde hedges that his personal guess is "probably not" malinvestment, but he treats the question as genuinely open.

"Is that enough? At what point does the extraction from these large two companies actually matter? Do we reach a point where everything we've done up until now is completely rational and everything we do into the future is complete malinvestment?"
— Wilde on the unresolved marginal-demand question

Open source, ad tech, and the compute-efficiency curve

Wilde is aggressively bullish on the acceleration of open-source model adoption as compute efficiency improves. As models become cheaper to run, more of the intelligence stack shifts to whichever open-weight release is closest to the frontier at a given moment, and the frontier labs' pricing power on their API business degrades accordingly.

The parallel bullish story is what this does to advertising. Wilde observes that the median ad on YouTube and TikTok is already effectively AI-generated, and that trend is one-way.

The next step is per-viewer generation — ads produced in real time for a specific person's inferred preferences. The economic upending of the advertising sector is, in his view, one of the most underappreciated consequences of near-zero-cost real-time content generation. It is a first-order thesis for anyone underwriting the ad-supported media stack.

The national-security case for an infinite demand ceiling

For an individual consumer, the marginal utility of more intelligence probably does hit a Pareto-shaped ceiling. But for a nation-state — specifically, in the current U.S.-versus-China dynamic — there is no upper bound. Falling behind on frontier AI produces a structural disadvantage that no rational government can accept.

He argued that both the railroad and telecom booms had significant national-security dimensions at the time, and that current hyperscaler capital-allocation decisions are, at their core, not driven by geopolitics but by return math.

The 20% hurdle rate

The cleanest piece of allocation math in the entire conversation is Wilde's statement about hyperscaler hurdle rates. The lowest all-in return he has seen any of Google, Meta, Microsoft, OpenAI, or Anthropic model on a data-center build is 20%.

As long as that number holds — and as long as leadership at every hyperscaler shares the view that there is no cap to intelligence demand — the industry will keep allocating capital to compute rather than other uses.

"The lowest number I've seen any of those companies model is 20% all-in cost. As long as we're living in that world, resources are just going to continue to be thrown at this thing forever. Why would you buy bonds?" — Wilde on the hyperscaler hurdle rate and its consequence for capital markets

If every marginal dollar of hyperscaler capital can earn a modeled 20% by building compute, the cost of capital for every other use in the economy has to rise to compete.

That is Wilde's explanation for why yields on Bitcoin-miner and neo-cloud debt at 7% to Tier-1 hyperscaler tenants look attractive relative to Treasuries — the data-center opportunity set is pulling the entire term structure of risk-adjusted returns higher. It is also a reasonable explanation for why the AI capex cycle has been so much more

capital-durable than skeptics predicted; the alternative uses of that capital simply do not clear the same hurdle.

Three Main Takeaways From The Podcast

1. **The bottleneck is moving.** In 2023 it was compute. In 2024 it was memory. In 2025 it was power. As of late 2026, on Wilde's ground-level read, sensor manufacturing and the human-data pipeline that feeds physical AI are the bottleneck. That progression is a useful map for anyone trying to figure out where the next constraint — and therefore the next set of pricing-power beneficiaries — will show up.
2. **Less work is a feature, not a bug.** Wilde's single most useful contribution to public-market investors may be his demonstration that hours worked and returns generated are almost entirely uncorrelated in this discipline. The best use of an investor's time is to cultivate a network of specialized nerds who bring pre-filtered ideas, develop a preconceived understanding of a small number of themes, and sit on cash the rest of the time. In his experience, every additional earnings release read is a rounding error against actual returns.
3. **Individual leverage keeps compounding.** The recurring counterpoint to any concentration-of-power worry about the frontier labs is that the cost to build almost anything on top of them keeps dropping. A system that took \$200,000 to build last year now takes \$20,000 with three weeks of work. That is the environment in which Inheritance AI itself was founded, and it is the environment that will produce the next generation of specialized physical-AI applications, mining-automation startups, and single-purpose robots that Wilde is most excited about.

I hope you guys enjoyed this conversation!