The Final Offshoring

Jacob Rintamaki, January 2026

Thank you to Verda Korzeniewski, Pranav Myana, Shreyan Jain, Zebo Furqatzoda, Michael Retchin, Sholto Douglas, Hiranmay Darshane, Teddy Ganea, Konstantine Buhler, Alexander Kesin, Tyler Cosgrove, based16z, Phillip Bogdanov, Alana Goyal, Michael Ferro, Stash Pomichter, Abhi Desai, Mohit Agarawal, Jack Altman, Bala Chandrasekaran, Sofia Puccini, Jimmy Apples, and many others.

Table of Contents

VII. Utopia's Paradox
He had been running all his life, from Fresno to MIT to one of the world's leading robotics companies. But what will this star engineer do when he's no longer needed to build the future?

VIII. Unit Economics
An adjustable financial model for the Minimum Viable Robot (MVR). Use it to build your intuitions, test out assumptions, and come to your own conclusions.

I. Why Robotics Will Work

If you try to look to the past to understand the future of robotics, then the future is bleak.

Dancing humanoids may crowd the trade show halls of Shenzhen but factory floors are still conspicuously empty. Most robot companies that have not gone bankrupt have stayed confined to household appliances[1], narrow industrial automation, or "boxes with wheels," whether iRobot with its Roomba, Locus with their warehouse AMRs, or, more recently, Waymo with its self-driving cars.

Thus, why should the future be any different? Why should one expect a sudden, dramatic wave of robotics working not just in the coming decade, but the coming handful of years? Why should the curse of Moravec's Paradox suddenly break?

The standard answer a savvy technologist would give is that increasingly capable AI video and world models will serve as a "base," providing real-world understanding, while deployments, whether through teleoperation, data gloves, or egocentric capture[2], will generate an additional data flywheel. This has already led to interesting emergent behaviors: absorbing egocentric data, tactile sensing, and generalization across environments. And we're about to scale everything up by 100x. Long robotics. Things will be big soon.

I think this is mostly correct, but let me add some nuance around both why to be bullish and two of the challenges that robotics faces today.

* * *

The Nuanced Bull Case

First, robot deployment is valuable even beyond just being a data flywheel for robotics models. For one, it seems like real-world data on your specific robot embodiment is the best kind of data to collect. But, in addition to that, the increased ability of Vision-Language-Action (VLA) models to "absorb" egocentric data is directly related to how much diverse robot data the model was trained on. So, every time we deploy a robot, it can now be "amplified" by egocentric data.[3] Furthermore, as we deploy more robots, the industry will gain morale, a lower cost of capital from creditors, and a better understanding of how to solve real problems for real customers, rather than just being stuck in a lab forever.

Second, I think most still don't realize how unoptimized everything is in robotics compared to what it could be. For the longest time, robotics' core problem was its lack of data, which needed ALOHA, UMI, and Diffusion Policy to break the chains. However, now that we are able to stably

train on and acquire increasingly large quantities of data, we can start changing many layers of the stack, from pretraining to teleoperation systems to even trying to hyperoptimize robotics via speedruns[4], similar to LLM speedruns such as NanoGPT.

This lack of optimization[5] is why there are so many different approaches to robotics at the technical level[6], from simulation-focused companies to those focused on diverse data to those focused on multiple real-world embodiments and techniques. It's also why the "robotics foundational model" layer is more contested than one would expect, which is why many are becoming "full-stack" competitors, as we'll later describe.

Third, robots do not need to be perfect initially to have a big enough economic impact to keep the "hype flywheel" spinning. As we'll discuss in the markets section, one framing of general-purpose robotics that I haven't seen much of isn't that we now have a robot that can do anything, but rather we have a robot which can quickly, cheaply, and easily be made to do one thing very well. If you think about the size of the TAM in jobs involving narrow, repetitive labor then, well, there you go. There's your flywheel.

⋆ ⋆ ⋆

Challenges

First, "evals"[7] for robots are very primitive compared to LLMs. As Kyle Vedder of Physical Intelligence put it, "LLMs differ from robotics…LLMs are able to be rolled out an unlimited number of times from the identical state s. [Robots, by comparison, cannot be rolled out an unlimited number of times from the identical state s.]" Furthermore, unlike the Kaplan or Chinchilla scaling laws[8], robotics has to deal with different embodiments, environments, tasks, and other small nuances which make extracting a single "log-linear curve of destiny" quite difficult. However, I believe much of this is also due to the relatively primitive state of robotics and world models, at least compared to where it will be in the near future. Furthermore, LLM evals will also need to be updated to account for progress. For example, METR's long-horizon tasks evals might need to switch to measuring "human programming uplift,"[9] as ex-OpenAI researcher Daniel Kokotajlo proposes in the depths of the AI Futures Project's new forecasting model. Similarly, robotics will also need evals that measure long-horizon capabilities[10], in addition to generalization, which is described as being both "faked" in many companies' demos and the source of controversy over whether "fine-tuning" is general enough.[11]

Second, I believe that the speed at which robots do tasks is asymmetrically important for the rest of the world. While an astute observer would note that the core problem with robotics has always been generalization, not speed[12], it does seem "quite lame" to have all robot demonstrations on 4x or 10x speed. However, we've already advanced far beyond naive speedups into SpeedTuning,

SpeedAug, advantage-conditioned RL[13], and training on "real-speed" data[14] rather than slow teleoperated data. I suspect a non-trivial portion of many people's skepticism will be washed away as "1x speed" demonstrations are increasingly shown; it certainly has for me.

* * *

Notes

1. Is a dishwasher a robot? Personally, I don't think so, but enough academic talks at Toyota Research Institute use it as a joke opening that it might as well be.
2. The most primitive version of this being an iPhone strapped to someone's head, recording their actions.
3. Amplification $\neq$ replacement. Keep that in mind.
4. I really think people did not pay enough attention to the top competitors here. There is a lot of alpha to be found from these types of challenges/communities, regardless of what field they're in. I'm very excited to see what crazy architectures, data mixes, and "clever hacks" we'll be using in 2030. It'll be astounding.
5. Which is exciting! There are so many avenues to pursue.
6. As compared to LLM companies, which tend to have more or less the same pre, mid, and post-training recipes these days. Yes, this is a gross exaggeration, but on a relative basis versus robotics companies, this is true.
7. Evaluations. Here is an interesting paper Anthropic did on how we can take a statistical approach to our evals.
8. And I want to be very careful referring to scaling laws here, as Aidan Clark of OpenAI rightfully pointed out that there is a lot of nuance when it comes to describing these. Treat this all as a starting point, rather than a finish, for robotics scaling laws.
9. Essentially, could I finish a similarly sized program 10x faster now than I could before? Also, Flowers for Algernon should've talked about Opus 4.5 in Claude Code.
10. While this demo ran for 13 hours before the operators decided to shut it off, as demos expand in length, it will be increasingly challenging to evaluate these policies. This is, in a sense, another argument for deployment: it allows one to run A/B testing at scale for new policies, but if world models can progress fast enough, then those could be a compelling alternative.

11. I think it's bullish for Physical Intelligence that they have this recipe for completing tasks, but it still requires fine-tuning rather than being done naively from the base model. I suspect scale and new techniques will fix this, but it is a question that a savvy robotics bear would ask.

12. As industrial robots can operate quite fast, although they do need safety guardrails.

13. Which did have to rely on human DAgger data, although it is promising.

14. Such as UMI-based, egocentric, or video/world model data.

II. Who Will Make The Robots?

After the technical section, I think many people would have questions about both the markets and the supply chains for robots. However, the more pressing questions would be about supply chains. The reason for this is that even though there will certainly be a lot of nuance and better versus worse markets for robots, many people will think the following:

> "Wow: home robots will be big. Wow: manufacturing will be big. Wow: many people work physical jobs. Thus, there will be big companies. We all know that. So, how do we not lose robotics to China? (Since it seems like we're on track to lose to China.) Not only are they good at making things™ and America is bad at making things™, but they also have dominance in rare earths processing and actuator manufacturing, which we don't.[1] Yes, Tesla is using the automotive supply chain to try and vertically integrate the production of Optimus, but will that be enough?"

Here is my response to add some nuance to this debate, first covering the high-level, then the low-level details.

⋆ ⋆ ⋆

To start, there is an assumption that America has to do everything itself. That's not true, particularly at the start, given our currently weak robotics hardware ecosystem. Vietnam is the obvious "China plus one"[2] but Malaysia[3], Taiwan[4], South Korea[5], and Japan[6] all offer interesting opportunities.

However, underestimating China would be a very grave mistake, and I don't want to be perceived as flippant here. To start, even large infusions of capital cannot easily replicate the highly competitive, deeply networked ecosystems[7] which make up the Chinese robotics and broader manufacturing ecosystem. Secondly, while there have been backdoors which have been found in companies such as Unitree, many are still using these robots because there are no realistic alternatives at the moment to use. Granted, Unitree is mostly used for entertainment purposes rather than industrial ones, but it should be concerning that there are hardly any "buy now" buttons for most non-Chinese robotics companies.

⋆ ⋆ ⋆

Rare Earths

Furthermore, if you trace the robotics stack from the top down, from the finished humanoid to the raw inputs, you will notice that minerals are upstream of almost everything. The robot needs actuators. Actuators need motors. Motors need magnets. And magnets need rare earths.

China controls roughly 85% of global rare earth processing and 90% of rare earth magnet production, and has already begun to impose export controls. While, yes, the geographical distribution of rare earths is fairly abundant, the processing capabilities are not, and one can just look to Standard Oil to understand how important monopolizing processing capabilities before taking over the entire "stack" can be for a company or nation.

However, there is an interesting solution to this problem: repurpose the Oil & Gas industry to help solve rare earths processing. It's already being discussed in the context of finding deposits of critical minerals, but these companies could be valuable for the actual processing itself. The chemistry (not to be flippant) is just liquid-liquid separations and solvent extraction, both of which the oil and gas industry does at massive scale every single day. Exxon separates hydrocarbon fractions at 99.9% purity while processing millions of barrels.

Granted, the specific separation chemistries for rare earths are harder, but the operational expertise directly transfers. Who else has been running large-scale continuous chemical processes, handling hazardous materials, managing regulatory complexity, and operating 24/7 with minimal down times for many decades?

So why did oil and gas companies not do this already? Because there was no reason to. Chinese subsidies made it uneconomical for everyone else, and no strategic imperative existed to change that. But the incentives have flipped. The Trump administration is talking about reshoring American manufacturing with robotics. JPMorgan Chase has committed $1.5 trillion to a Security and Resilience program with rare earths as a focus. MP Materials saw the Department of Defense take a $400 million equity stake, making the US government its largest shareholder, plus a 10-year price floor commitment and a $150 million loan to expand heavy rare earth separation capacity. Lynas achieved the first commercial separation of heavy rare earths outside China in decades.

All of this seems like good news to me. Of course, there is concern that the solutions will not ramp in time, but there are a variety of "bridge" solutions that we could utilize[8] to help smooth things out.

⋆ ⋆ ⋆

Actuators

The real worry, to me, is more so in actuators, which is really a problem of precision reducers.

A precision reducer is a mechanical device that sits between a motor and a joint. The motor spins super fast but with a low torque, which the reducer then converts into a slow, precise motion. The two main types are harmonic drives, which use a flexible gear that deforms inside a rigid outer gear, and RV reducers, which use a planetary gear system. But both require manufacturing precision measured in microns.

And that's where the trouble starts.[9]

The Japanese conglomerates Harmonic Drive and Nabtesco dominate this market. Nabtesco alone holds 60% of the global market for RV reducers. These components are difficult to replace or scale because certain failure modes only emerge after thousands of hours of operation. Tribology, the study of friction and wear between interacting surfaces, is devilishly complex.[10] The knowledge is empirical and accumulates slowly.

You cannot simply buy your way to competence.

Or can you?

The implicit assumption is that we need Harmonic Drive quality for every application. We do not; most robotic tasks are not brain surgery. The robot folding your laundry does not need the same precision as the robot performing your appendectomy.

First, we can use cheaper, lower-precision components and replace them more frequently. If the robot is generating enough revenue, swapping out a reducer once a year is not a dealbreaker.

Second, software can compensate for hardware imprecision. This is the whole point of learning-based approaches. You adapt to whatever embodiment you have. The ALOHA system runs on cheap motors with 3D-printed gears. Physical Intelligence's recent work shows you can do real manipulation tasks with simple hardware. Most industrial robot failures are integration and software issues, not precision reducer wear. Even FANUC replaces components on a regular schedule.

Third, Chinese manufacturers are catching up. Green Harmonic, Leaderdrive, Zhongda Leader, Shuanghuan Transmission. They are shipping reducers at 80 to 90 percent of Japanese quality for 30 to 40 percent of the price. This matters not because American companies will necessarily use Chinese reducers but rather because of the "four-minute mile effect."

Once the barrier is broken, everyone knows it can be done.

China has proven that cheap, good-enough reducers can be manufactured at scale. Which means

factories in Vietnam and Malaysia can follow, as the ecosystems in Bac Ninh and Penang are deeply underrated.[11]

Fourth, qualification cycles compress when you run them in parallel with deployment. Every hour of field operation is a new hour of qualification data. The multi-year timelines assume you qualify parts before deployment. If you qualify during deployment, accepting higher early replacement rates in exchange for faster iteration, the math changes entirely.

* * *

The Cost of Robots

The implication from this is that robots will cost closer to an iPhone than to a car.

There are a few ways you can arrive at my conclusion. The first is simply just by looking at the size of a robot (which could be just arms, or just arms and a wheeled base) versus a car. Why would a normal robot, which is made out of much less raw materials and does not need to be safety qualified the same way a vehicle which could move over 100 mph is, cost the same from a hardware OEM standpoint?[12]

The second is that "robots making robots," which is often dismissed as a "sci-fi" or "deus ex machina" explanation for why the cost of producing robots could collapse, is not as science fiction of a concept as it sounds. FANUC has been doing this, albeit in a much, much narrower way, since 2001. As robots get more general, we can just gradually expand what FANUC has been doing for 25 years now.

However, regardless of how you reach this conclusion, it's quite important because it will shrink the payback time for robots down by quite a bit. This, after default risk, is a major thing that creditors look for, and if an increasing amount of credit can be poured into robotics, then it could kick off many, many flywheels for the industry.

* * *

Notes

1. Furthermore, there might be some small odd bottlenecks throughout the supply chain that they'd be better positioned to fix, but that seems unlikely to matter all that much. While certain robotics markets are not similar to the automotive industry, much of the robotics supply chain

will approximate the automotive industry, which produces random odd components in mass volumes.

2. It has a low cost of labor, an educated workforce and an eager government. Yet, relationships there are getting competitive fast: everyone from Samsung to large portions of SpaceX's Starlink supply chain is beginning to show up.

3. Penang's semiconductor and electronics manufacturing ecosystem provides deep bench strength in precision assembly.

4. Underrated. Everyone avoids it because of invasion risk, which (1) means the people willing to tolerate that risk get opportunities others won't, (2) recent geopolitical moves from the US indicate that moves for "defending Taiwan" may be on the table now, and (3) there is a massive market in servicing the AI datacenter buildout happening in Taiwan right now. If your customer and your manufacturer are the same entity, your feedback loops get very fast.

5. Their shipbuilding industry is well-known (although it's less well known that Hyundai owns Boston Dynamics) but they have the highest robot density in the world. Additionally, Samsung and SK Hynix are both deeply involved in the AI data center buildout.

6. FANUC, Mitsubishi, Honda, Toyota (and its famous research institute, TRI), Tokyo Electron, Sony, Softbank…many things to consider.

7. Shenzhen, most famously.

8. Recycling or reclaiming rare earth magnets, utilizing the existing inventory buffer for robotics, and cutting unnecessary site permitting regulations.

9. A micron is one-millionth of a meter. A human hair is about 70 microns. These gears need tolerances of 1 to 5 microns across surfaces that will rotate millions of times under load. This is not something you can achieve with a CNC machine ordered off Alibaba.

10. Tribology is the devil of mechanical engineering. How surfaces wear depends on hundreds of variables that interact nonlinearly, down to the molecular level. You cannot simulate your way to understanding. You have to run the parts until they fail, then figure out why, then try again.

11. Samsung alone employs 40,000 to 50,000 people in Bac Ninh. The skills for precision manufacturing already exist, they just have not been pointed at reducers yet.

12. Yes, precision hardware is expensive, but you can already get very cheap, non-industrially rated arms for a couple hundred dollars thanks to HuggingFace.

III. Follow The Money

> "You Americans measure profitability by a ratio. There's a problem with that. No banks accept deposits denominated in ratios. The way we measure profitability is in 'tons of money.' "
> —Morris Chang, founder of TSMC

If you're like me, you just read the previous section and are now quite bullish on robotics happening. However, before we get into the much more speculative side of things, I think that it is important to understand what the markets for robotics will look like.[1]

I believe this is important because there are likely many interesting information asymmetries to be found here with careful analysis.

A high quality view of the markets in robotics can be found in "The 2026 AI Forecast with Sarah & Elad". Their viewpoints can be roughly blended in with what is my understanding of smart venture investors' perspectives on robotics as the following: Incumbents are favored because robotics will be capital-intensive and require hardware/supply chain expertise. Yes, Tesla, SpaceX, and Anduril were able to pull it off, but those were abnormal startups. Granted, there will still be startups that do well here, but incumbents like Tesla and the Chinese companies will be very strong. Additionally, there's a lot of uncertainty around how the market will shape up, but it seems like there will be markets around data collection[2], horizontal software, enterprise robotics, home robots, hardware OEMs, and robotics model providers.

Personally, I think this seems to be mostly correct, but there are a lot of details we could layer in here. We'll first address some of the "macro variables" that could affect these markets, and then address enterprise robots, hardware OEMs, robotics model providers, data collection/horizontal software, and home robots. "Robot specific" markets which will emerge over time will be addressed in later sections.

⋆ ⋆ ⋆

Macro Variables

To start, let's address those "macro variables."

Firstly, the funding environment for robotics companies will initially come from venture capital. Deals will be done initially more so by the multi-stage funds than preseed/seed funds[3], and there will be dramatic valuation step-ups after robotics companies win enterprise contracts, similar to defense companies. These "stepped up" companies will be especially interested in strategic VC

arms, as they will be ways to secure further large enterprise contracts.[4]

Secondly, credit will play a larger role in robotics than they do for traditional tech startups because robotics has an industrials component which B2B SaaS doesn't. So, the key variable for companies to focus on will be "payback times,"[5] in addition to "what is the equity multiple over ~3-5 years once you leverage your asset pools?" RaaS-based financing from domain experts will also become increasingly valuable compared to taking on traditional private credit or venture debt.[6] But, interestingly enough, private equity (such as EQT and 1X's deal) will have a mutually symbiotic relationship with robotics companies, not just because many PE funds also run massive private credit businesses and venture arms, but because robotics can improve the performance of their portfolio companies.

Thirdly, Asia will not just manufacture robots; they will also buy robots. Many might dismiss Asia as a market because of their lower cost of labor[7], but rising elderly populations, AI datacenter booms causing demand increases[8], rising labor costs, population decline, fertility decline, government policies, and a high density of manufacturing robot deployment all point towards the markets here being unusually interesting.

And then there is the policy tailwind, because The One Big Beautiful Bill Act (OBBA) includes 100% bonus depreciation for qualifying equipment.

It means that you can depreciate the entire cost of a relevant piece of equipment in year one instead of spreading it over five to seven years.

Both robots and AI infrastructure qualify.

And there is no cap.[9]

Off to the races!

* * *

Enterprise Robots

Given the above, it seems like enterprise robots will be adopted surprisingly quickly.[10] Similar to the wave of "AI Application" companies that exist for LLMs, I believe we'll see the rise of the "Minimum Viable Robot" (MVR) company.

A wheeled (or no mobility)-based robot with parallel-jaw grippers running a frontier open-source VLA post-trained on custom data is rapidly deployed into environments with "teleoperation as a bridge[11] to full autonomy," which would collapse both the OpEx and thus the payback times for

the system. (A more rigorous treatment of the unit economics here, particularly around speed and teleoperation "intensity," can be found in the "MVR Unit Economics calculator" in Section VIII.) The strategy for these companies then, given that reducing payback time may be All You Need, is to deploy into large enterprise customers as aggressively as possible to start building moats that their larger video/world-model focused competitors[12] still find difficult to match, similar to the best AI Application layer companies. This is also similar to the strategy which have made the "systems integrator" companies, a la ABB, Siemens, FANUC, KUKA, and so on, remarkably durable: deploy, deploy, deploy into customers, and become their go-to person for fixing messy problems.

Within the "enterprise," of course, there are many potential markets to go after. As just a shortlist: Logistics[13], Healthcare[14], Retail[15], Agriculture[16], Cleaning[17], Hospitality[18], Manufacturing[19], and Restaurants[20].

However, the most interesting and underdiscussed set of markets for enterprise "MVR" companies would be anything related to the AI datacenter buildout.

⋆ ⋆ ⋆

The AI Datacenter Buildout

The US economy has become a levered bet on AI infrastructure, and that bet shows no signs of stopping.

People compare this to ZIRP, the zero-interest rate era, when money was free and companies could lose billions indefinitely. The comparison is wrong, though, as ZIRP was a macroeconomic phenomenon. Money was cheap, so companies spent it, but when the money became expensive again, the spending stopped. The underlying demand was never real; it was conjured by financial conditions and vanished when those conditions changed.

AI is different, because AI is a technological, not a financial, phenomenon. The spending continues not because money is cheap but because demand is real and growing faster than supply. The hyperscalers are not building datacenters because they have money to burn. They are building datacenters because every GPU they can bring online is immediately utilized. In 2000, telecom companies laid fiber expecting voice traffic to grow exponentially. It did not. The dark fiber sat unused for years. The AI buildout is different: the demand is arriving before the infrastructure is even built.

Only 13 gigawatts out of 125 gigawatts of existing datacenter capacity has been converted to GPU computing. Thirteen gigawatts is the output of thirteen nuclear reactors or the total electricity consumption of the Netherlands.[21] And this is only what has been converted. The planned capacity

is five to ten times larger. Consumer demand, enterprise adoption, sovereign AI initiatives, video model training, and foundation model scaling all will compound on itself. Are you sure that you want to bet against that?

People always ask what could stop the buildout. The objections fall into three categories.

Objection 1: Model Progress

The first objection is about model progress. What if the next generation of models shows diminishing returns? What if GPT-6 is only twice as capable as GPT-5 instead of ten times? What if boards start asking hard questions about ROI? Of course not! That'd be crazy talk.

The answer is that continued spending does not require continued breakthroughs. For the buildout to stop, revenue would need to stop as well. And for the revenue to stop, the models would need to stop improving entirely. A diminishing return is still a positive return. AWS continued building datacenters even as growth slowed from 40% to 15% annually, because the absolute numbers remained enormous. Even in a slowdown scenario, inference demand will grow rapidly for years. The strongest version of this objection is the "good enough" plateau: what if GPT-4 class models satisfy 90% of enterprise use cases and there's no economic reason to build bigger? This is possible, but inference demand alone would still require massive infrastructure growth even if training stopped entirely.

Some claim that on-device inference will absorb demand and chill the datacenter buildout. This misunderstands the market. On-device inference and datacenter inference serve different workloads. On-device inference will summarize your iMessage history. Datacenter inference will automate white-collar work. The compute requirements will also differ by orders of magnitude.[22] These are complementary markets.

Objection 2: Infrastructure Bottlenecks

The second objection is about infrastructure bottlenecks. The development cycle for GPUs is fast, but the deployment cycle is slow. If GPUs sit idle between utility failures, equipment delays, and labor shortages, then training capacity (which is a key part of ROI calculations) will suffer.

This is a problem, but the hyperscalers can brute-force their way through. Alphabet just bought a Texas energy company, Intersect, for $4.75 billion.[23] Amazon has signed nuclear partnerships for 1,920 megawatts of carbon-free nuclear power through 2042. Microsoft signed what Brookfield called "almost eight times larger than the largest single corporate PPA ever signed," totaling 10.5 gigawatts. Meta released an RFP targeting 1-4 gigawatts of new nuclear capacity. The pattern is consistent across all major players. They are not waiting for utilities to solve their problems. They are becoming utilities.

Stargate has also been up and running since September. Granted, in initial operations, not at full capacity, but the capital is committed and construction is underway.[24] OpenAI announced that

combined capacity from their new sites brings Stargate to nearly 7 gigawatts of planned capacity and over $400 billion in investment over the next three years. Additionally, the UAE has another benefit for its partners: an infinite money printer called "oil and gas reserves" that it plans to convert into AI infrastructure. The UAE has cheap energy, no permitting delays, and a $1.5 trillion sovereign wealth fund. They are building AI infrastructure faster than the West can approve environmental impact statements. UAE capacity is a hedge for the hyperscalers and demonstrates that the buildout will happen somewhere.

Objection 3: Geopolitical Risk

The third objection is about the serious geopolitical risk facing AI.

Everyone in this industry knows that China is a geopolitical threat to both the sole maker of advanced chips and 90% of global AI server manufacturing. But this is not new information, as TSMC is building in Arizona and Japan, Samsung is expanding in Texas, and the hyperscalers are diversifying.

I will not pretend this risk is small, however. A Taiwan crisis severe enough to halt chip production would halt the global economy and supply chains would shatter. Additionally, while the diversification is real, it is slow: TSMC Arizona has already experienced significant delays, with the facility now expected to be fully operational in 2027 or 2028 rather than the original 2026 timeline. Samsung Texas has also faced similar setbacks.

Furthermore, if Taiwan falls, then Japan and South Korea will also face cascading consequences due to their reliance on the South China Sea for shipping. If you try to use either country as a "hedge," you should receive valid criticism along these lines.

However, I'm still personally optimistic about Taiwan, especially considering that, for better or worse, the US government has recently shown that it is willing to take unconventional measures. Traditional geopolitical analyses will likely not apply well here.

* * *

Robotics Markets in the AI Datacenter Buildout

The AI datacenter buildout will create four distinct markets for robotics, each with different characteristics and different levels of difficulty. The smart strategy here will be to start where the problems are tractable for MVR systems and then concentrically expand one's TAM outwards.

Datacenter Maintenance

Datacenter maintenance is the easiest entry point for MVR systems.

Why?

First, the environment is climate-controlled and standardized.

Second, the tasks (inspection, cable management, hardware swaps) are repetitive.

Third, the buyers are sophisticated and have ample budgets.

Fourth, the failure modes here are non-catastrophic.

Fifth, automated inspection with cameras on mobile bases is already proven technology. Managing cables, maintaining racks, handling logistics, and moving hardware between zones do not require a superhumanly capable robot; you just need something reliable that operates around the clock in harsh conditions. Granted, the reliability requirements here are real, particularly for the hyperscalers, but they are achievable with today's technology. And, unlike self-driving, if a robot drops a cable, it won't (well, at least it shouldn't) accidentally kill someone.

Supply Chain & Manufacturing

The supply chain and manufacturing markets for the AI datacenter buildout is next up on the difficulty curve, because datacenters require enormous quantities of equipment: IT racks, servers, networking devices, power distribution units, and cooling systems. Even small datacenters contain hundreds of thousands of electrical components, which is part of why Foxconn now makes more money on AI servers than iPhones and NVIDIA's networking business is becoming its fastest-growing segment. Additionally, NVIDIA is shifting to 800 VDC power architecture starting in 2027, which means over 13 gigawatts of installed GPU clusters will need to be retrofitted.

The interesting thing about this market is that it consists of repeatable work in fixed locations. Additionally, the entire global electronics supply chain is reorienting around AI infrastructure, since the same robots that assemble servers can assemble electronics for phones, cars, appliances, whatever you want! It's all the same companies doing the work anyway: Foxconn, Quanta, Wistron, Pegatron, etc.

Datacenter Construction

Datacenter construction is perhaps an even larger market, but it is also the most complex for MVR systems to attack. The key insight here is to separate post-shell work from pre-shell work.

Post-shell work is what happens after the building envelope is complete: installing HVAC, running cables, mounting equipment, and setting up networking. The environment is controlled and the tasks are standardized here, so a robot installing prefabricated server racks in a climate-controlled shell is performing work similar to that already performed by robots in warehouses. A hyperscaler building their tenth datacenter has figured out exactly how they want it done. The specs are

documented, so it's closer to manufacturing than construction. Here are some rough estimates for the cost breakdowns: post-shell work is probably 30-40% of total datacenter construction cost. Pre-shell is 40-50%. The remainder is land, permits, and soft costs.

Pre-shell work is harder, though. You'll need to deal with structural work, electrical, plumbing, pouring foundations, erecting steel, running conduit through concrete, bespoke designs, constantly changing conditions, heavy equipment coordination, permitting, and inspection requirements. Pre-shell construction is where robotics dreams go to die. Boston Dynamics has been demoing construction robots for a decade. None are deployed at scale. The chaos of the job sites is non-trivial to deal with.

Unlike Katerra, the prefabrication startup that failed because it tried to impose standardization on an industry that was not ready for it, the post-shell opportunity is real. Katerra believed the same thing I'm arguing: that volume forces standardization. They failed because the volume wasn't there yet and they burned cash trying to create it. The difference with datacenters is that the volume already exists and is growing.

Thus, the path forward here is clear: start with post-shell, expand into pre-shell as capability matures, and let the technology pull you into adjacent markets as it improves.

The Flywheel

All of this adds up to something robotics companies have never had before: a market that is simultaneously large, fast-growing, high-margin, and served by sophisticated buyers who will pay premium prices for solutions that work.

The flywheel is elegant: Robots will build AI infrastructure. AI infrastructure will train better models. Better models will make robots smarter. Smarter robots will then build more AI infrastructure, as each turn of the wheel accelerates the next, faster and faster.[25]

⋆ ⋆ ⋆

Hardware OEMs

Compared to the enterprise market, the hardware OEMs are somewhat easy to predict. They'll mostly be automotive OEMs in Asia that are repurposed to serve the robotics industry, and there will be a niche of very high quality robots, similar to the supercar market, that will exist.

⋆ ⋆ ⋆

Model Companies

The model companies are somewhat odd, though. It seems like it will skew heavily towards incumbents due to the capital intensiveness of training large video models, and if we "pattern-match"[26] there should be 1-2 closed-source startups doing this, along with a handful of open source model providers, likely originating in China. As mentioned before, going full-stack into enterprise deployment is likely the best move for existing "model companies."

⋆ ⋆ ⋆

Data Collection & Horizontal Software

However, when it comes to the "Scale AI for robotics" discourse, my opinion is that it seems somewhat overblown. In some sense, "the Scale AI of robotics is Scale AI,"[27] but there will be a handful of companies which emerge. Providing value-added services and cornering experts, embodiment specific, or regulated will be ways to differentiate here, in addition to operational expertise. However, I find it hard to believe that this by itself will be an exceptional market. Layering in horizontal software, such as DevOps, fleet management[28], simulation/RL tooling, or a "software/firmware OS"[29] could be quite valuable. In some sense, it reminds me of the GE Aerospace strategy: sell engines at cost to lock customers into decades of regulated high gross margin services.

⋆ ⋆ ⋆

Home Robots

Oddly enough, in the long run, Home Robots may have similar models to the GE Aerospace strategy, although without the regulatory arbitrage part. Of course, it should be noted that home robots are overrated in the short term. The issues of distributed repair and maintenance versus enterprise deployments, safety concerns, high upfront costs that most households can't afford, data privacy concerns, unrealistically high expectations (plus "uncanny valley" effects) caused by the human form, and the challenging "out-of-distribution" characteristics of many homes make it a difficult product to sell currently. However, in the long term, home robots will not resemble the automotive industry; they'll feel a lot more like the iPhone. The core use case for home robots, chores, is just the bare minimum of what they could do. If you dropped the marginal cost of home labor to zero, then the ability to provide a five star service hotel to any customer would be within reach, and that's without being creative.

* * *

Notes

1. Or, more glibly, to "follow the money."
2. A "Scale AI" for robotics.
3. Which need to hit certain ownership targets to make their economics work.
4. GS Futures, Toyota Ventures, and Prologis Ventures are examples of these "strategics" VC arms.
5. An interesting thing about robotics is that you should be able to still get quite good IRRs, even after you haircut for customer default risk, utilization dips, servicing/field ops, and your real cost of capital (using 10–12% senior, 18–25% equity).
6. Also, if you're worried about being exposed to volatile rates because of taking on a large amount of debt, rate collars are preferred over swaps. Companies should want to allocate margin to R&D and go-to-market rather than being too "gigabrained" on hedging.
7. So, thus, lower revenue ceilings for robotics companies.
8. Taiwan alone has had multiple 8% GDP quarters in 2025.
9. A company buying $100 million in robots gets $21 million in immediate tax savings, i.e. the effective cost drops 21% on day one. State tax treatment varies significantly, however, as California, for instance, doesn't conform to federal bonus depreciation. Thus, companies deploying in multiple states should model this carefully. However, let me add in one clarification: if your business had a contract to acquire property before January 19, 2025, even if the actual acquisition happens afterwards, the property will not qualify for 100% bonus depreciation.
10. This is a historical trend that can be borrowed from the industrial revolution. Centralized crankshafts brought power to factories first, before decentralized electric motors made it more mainstream. It's one of the justifications for my take that we will see enterprise robotics at scale before home robotics.
11. If you can teleoperate it, you can learn it.
12. Mag 7, The LLM Labs, etc.
13. Distribution is painful so don't try and be a narrow solution, don't bet your company on a single "big name" pilot and don't try to compete on raw speed for picking/packing. The ecosystem is also very in-person, heavily conference-oriented, and different from the circles

that many west coast AI researchers hang out in. However, I expect there to be many massive companies in this space, just because of how many, how large, and how forward-thinking many players here are for robot adoption.

14. If you don't know what a nosocomial infection is, you should probably go look that up. This market (which we are excluding surgical robotics from, as despite Intuitive Surgical's remarkable $206 billion success, they are playing in a very different market) will be more challenging than many expect. While some of the largest hospitals, such as Mayo Clinic, are very tech-forward, the "long tail" of hospitals frequently have disconnects between upper management and the actual day to day employees. I expect that technology diffusion here will be slower than expected, although eventually it will be implemented due to the sheer cost advantages robots will provide.

15. Retail is a K-shaped market. The top brands and stores will be very compelling customers while everyone else is slowly decaying. Interestingly, retail as a market collapses into logistics, as the largest retail companies, depending on your definition, are Walmart and Amazon.

16. Everyone talks about strawberry picking robots, but I haven't seen many talk about meat-packing robots. Part of the beauty of general-purpose robots is that they can do things that traditional custom automation cannot, which includes being able to handle heterogeneous SKUs. Agricultural products often involve a lot of this.

17. This is where the "Uber for robots" or similar types of business models will be most likely to emerge. However, it will become somewhat "knife-fight-y," so I would worry about that for businesses that choose to compete here.

18. If you are able to lock in enterprise contracts, then this could be an interesting market, as it is the closest enterprise equivalent for home robotics companies. Similar to home robotics, the aesthetics for these robots will matter. Unlike home robotics, the working hours, competitiveness with cleaning staff, and the relative inability to grow a context on what the user likes or dislikes are hindrances.

19. The automotive industry (and the consumer electronics industry, but we're just about to get to that) will be the most interesting market here, not only because they've been incredibly interested in adopting robots, but also because there are very few companies which are sophisticated users of automation and want to place large orders for robots. Additionally, if you are working with defense customers, then be sure that your teleoperation systems use US workers for them, otherwise you could be violating data privacy/ITAR regulations.

20. The idea of "bowl robots," arms that scoop for restaurants such as Sweetgreen, Cava, and Chipotle, is already within the public consciousness. However, expanding this out to focus on high-quality chains with reputable brands (which Ghost Kitchens do not, so those should be avoided initially) will likely lead to at least one large, "Toast-esque" successful company being formed.

21. Of course, not all of that 125 GW can be economically retrofitted for GPU computing. Legacy facilities, wrong locations, inadequate power density. A more realistic conversion target might be 50-60 GW, which is still enormous.

22. A 2B-parameter model on an iPhone and a trillion-parameter model for complex reasoning requires different workloads.

23. Per the December 2025 announcement: "Alphabet today announced a definitive agreement to acquire Intersect, which provides data center and energy infrastructure solutions, for $4.75 billion in cash, plus the assumption of debt."

24. Granted, in initial operations, not at full capacity, but the capital is committed and construction is underway.

25. Something something Nick Land, Retrocausality, The Transcendental Object At The End Of Time. You know the deal.

26. Be very, very careful when doing this unless you think there will be a similar underlying causal mechanism. Which, in this case, there is, since it's model training.

27. Or, rather, the incumbent data labeling companies of the world will continue to be incumbents.

28. "Uber for robots" will be a thing, but it will take some time. Many people will make a killing here, and, as we'll later discuss, labor becoming a commodity will mean it will behave more and more like oil. So, if you want to build the "Uber for robots," figuring out how to efficiently transport large quantities of labor quickly will be very valuable for you. Additionally, it should be noted that Alphabet, through a subsidiary, owns the original makers and primary maintainers of ROS/ROS2. That will be highly relevant in the coming years.

29. The Microsoft for robotics opportunity, given that many want to use cheap Chinese hardware, or cheap hardware from many countries, but their underlying firmware/software is often not that great.

IV. The Future Is Not Fixed

Now, before we get into more speculative territory, I want to be clear about something: we still have agency over what kind of world we're building. The future isn't fixed, and you shouldn't act like it is.

I bring this up because of a debate that's been raging in AI's circles recently. The gist is podcaster Dwarkesh Patel and Stanford economist Philip Trammell recently argued that, in a world where AI and robotics make capital a perfect substitute for labor, inequality will skyrocket and traditional policy tools won't be able to easily fix it.[1]

Unsurprisingly, this caused a lot of discourse.[2]

Here are my takeaways from all of this:

First, inequality will matter more in a world with abundant robotics and AI. In particular, the inequality of power, not just wealth, will matter, and the developing world could become unusually volatile.[3] Yet, if we're not in a "crazy world" where the rules of economics are violated, then these problems can be approached with traditional tools: public policy, differential technology development, and the like.

Second, if we are in a crazy world where capital perfectly substitutes for labor, then assuming traditional tools will still work is delusional. When COVID hit, America went from a normally functioning economy to shutting down the NBA and New York City within weeks, appropriating trillions for stimulus checks and PPE production. Now, imagine if something 10x/100x/1000x more important than COVID hit. Why would anyone assume that the reaction would or should be muted and typical?[4]

Personally, however, I believe we'll be in the "normal" world for the foreseeable future. This will still involve an extraordinary amount of change[5] occurring by historical standards, but not one where the fundamental rules of economics completely break down.[6] So with that: welcome, one and all, to the final offshoring.

* * *

Notes

1. See Patel and Trammell's essay, "Capital in the 22nd Century".
2. Here we go. Economists offered alternative models or suggestions to Patel and Trammell's ideas but didn't fully address their core claim that human preference for human labor would

dissolve over time. Some more "mainstream" folks rejected the claim that capital would ever be a perfect substitute for labor, despite still agreeing that inequality would be a big deal. Meanwhile, some in AI Safety rejected the premise that property rights would hold at all, as "why should superintelligent AI with a trillion robots respect my claim?" Others cited a Nature meta-analysis suggesting economic inequality doesn't equate to poor wellbeing, but proponents of "gradual disempowerment" found this logic naive, as it didn't solve the problem of there being an inequality of political power, among other things. The debate reached beyond the usual hyperonline AI circles, including even involving AOC's former chief of staff and Ben Thompson. There's more I haven't covered, but this is the 80/20 for the debate.

3. Which we'll cover in the following section.
4. Move 37 is as relevant as always.
5. Thank you Dean Ball and Seb Krier for being the Andy Masley of AI economics it seems.
6. We will be covering how to deal with a scenario like this though.

V. The Final Offshoring

If you could only say one word about what the final offshoring will bring, it would be this: deflation. If you make labor cheap and abundant, then, for industries which are not heavily regulated, it should deflate the costs of those industries.

The downstream implications of this are enormous.

We'll start with the Fed, which has the twin mandates of maintaining a 2% inflation rate and maximum employment. The maximum employment point, as tangentially mentioned by Ben Thompson, could be changed by a number of degrees (i.e. "Is podcasting a real job?"), and unemployment is such a catch all term that I don't think anything about that portion of the Fed would need to be changed because of the final offshoring. However, the 2% inflation target is something to watch out for. Typically, deflation is fine because it only affects one industry, so consumer spending can be reallocated to other industries. But if there's a continual broad deflation across the economy, then things could get weird. The Fed tends to fight deflation by cutting interest rates, but you cannot cut interest rates meaningfully below zero. Both Europe and Japan tried negative rates with mixed results at best. I'm optimistic though, as through a combination of evaluations and methodology updates, the Fed, the BLS, and a variety of other bodies could be better prepared to account for economic changes in an uncertain world.

However, assuming the Fed is fine (which it should be), the impact of deflation on ordinary Americans will be immense. We'll get into the (many) positive implications shortly, but the core challenge here will be that many jobs that are today worked by people will be impacted, because American jobs tend to have higher wages than jobs overseas. The people who are no longer working those jobs will rightfully have opinions about that!

Satya Nadella once said that "the true cost of energy is social permission." I think this will be the same for robotics, in that, if the American people are not given a compelling vision of the future that includes them, why would they go along with this?

We'll cover what happens internationally if we don't go forward with robotics (spoiler: it's bad), but I think there is a lot of worry from many Americans about everything turning "socialist" if large wealth, robot, land, VAT, or other taxes are immediately imposed on them. As conditions change, I think a variety of reasonable taxes will be passed under spending bills, but the first set of changes that could be done would be finding a way to work the leading robotics companies into a sort of American Sovereign Wealth Fund (SWF). Lutnick and Bessent have already discussed doing something like this, and it is quite popular in the American states (notably Alaska with its oil fund) and countries (Norway) where it's been tried. Who doesn't love free money (as well as orienting all of our incentives towards growth, rather than division)?

However, I will say, to temper the expectations of those who are very pro-SWF, that if we do a 5%

removal every year from the fund, spread across 300+ million people, then the fund will need to be enormous to cover a modern American's standard of living. Mass deflation will help with this, but I am skeptical that a pure SWF will be enough to cover something equivalent to our way of life today. Yet, in terms of political feasibility, starting with a SWF, along with keeping a close eye on the technical progress of robotics, should keep the US in a surprisingly stable place for quite some time.[1]

Speaking of stability, remember how I mentioned what would happen if we didn't go forward with robotics? It's time to talk about that now. If America forfeits being a global player in robotics, then, yes, China will likely surge to be the leader.[2] However, the more worrying perspective is that countries which heavily rely on remittance payments, such as Bangladesh, will suffer as robots replace workers in wealthy countries. Those displaced workers, now with a much weaker case for immigration, will go back to their home countries, which may also be implementing robotic automation, causing currency debasement and possible debt defaults. When that happens, the instability of the world will dramatically increase, particularly since many of the countries most vulnerable to remittance going away tend not to have pre-existing strong democratic institutions which can maintain stability in transitional periods. But, again, if the US can remain at the frontier in robotics, we can serve as a helping hand here.[3]

⋆ ⋆ ⋆

This leads us naturally into the more positive parts of deflation. In his essay, "Machines of Loving Grace," Dario Amodei[4] wrote about a concept he called the "marginal return to intelligence." The gist is that certain tasks, like scientific research, had high marginal returns to intelligence, and so would be unusually benefitted from advancements in AI. Similarly, I believe there are many tasks which have a high marginal return to labor.

To start, if your goal is pure wealth accumulation, then finding a way to get as involved with the AI datacenter buildout as possible via robotics should be your sole objective.[5]

However, if you have more varied goals, then the final offshoring will unlock many domains where applying massive labor to a problem finally makes economic sense.

The custom manufacturing of everything, in as high of a mix and as low of a volume as possible, will become feasible. One perspective is that this will look like a "super-Etsy" becoming feasible. This will all likely lead to something akin to an "American Shenzhen"[6] emerging, likely in Texas, as American natural resources, IP, and strong domestic markets fuel a manufacturing boom.

Similarly, there will be a "care boom," both in the form of custom care for humans, whether they're a child, elderly, disabled, or a disaster victim, or the environment, as potholes, trash-littered highways, and sullied forests get remediated.

Tangentially, as someone who tries to spend a lot of time in nature, when I'm not with the robots, it's always been astonishing to me how little of the world has still been explored in-depth today. Geologists, Archaeologists, Deep-sea exploration[7], and even exploring other astrophysical bodies could become open.

As we're doing all of this exploration of the natural world, though, robots will be assisting in explorations of both the artistic and scientific worlds. The abundance of custom manufacturing discussed earlier could also apply to artistic works at scale: custom paintings, jewelry, and sculptures all seem within reach. Similarly, custom labor-intensive agriculture, such as exotic flavorings, could be cultivated more cheaply[8], as well as labor-intensive science, like the combinatorial chemistry required for longer-lasting batteries or novel perfumes.[9]

* * *

Yet, I also believe that there's another framework we could be using to analyze what will happen to us. Classically, the four factors of production were land, labor, entrepreneurship, and capital. As labor becomes commoditized, each of the other three factors will change in strange and unexpected ways.

* * *

Land

Land will become more important in the world of the final offshoring. That should be obvious, but how so isn't.

I'll start with some predictions. I suspect that David Ricardo's law of rent will make a comeback. The land-owning class will return to prominence, and so, Georgism, the idea that we should tax land values rather than labor, will receive serious reconsideration.

From a consumer perspective, places that are pleasant to live become more valuable. Europe, with its walkable cities and moderate climate and historical depth, is more attractive than ever. In a world of synthetic abundance, authenticity appreciates.

From an enterprise perspective, energy-rich and lightly-regulated land becomes critical, leading to countries with abundant energy potential to become the new centers of industrial power.[10]

* * *

Capital

First, interest rates will become strange. What is the natural rate of interest when production constraints vanish? Interest rates theoretically reflect time preference and the marginal productivity of capital. If capital can produce almost anything with near-zero labor input, does capital become less valuable or more valuable as the binding constraint?

Second, trade theory will require revision. Ricardian comparative advantage holds that countries should specialize in what they produce relatively more efficiently. But if labor costs go to zero everywhere, relative labor productivity stops mattering. Energy costs, regulatory environments, raw material access, and proximity to customers will soon determine trade routes.

Third, if robots become a commoditized factor of production like oil, then we will need financial infrastructure to deal with this. Robot futures markets will emerge for the right to buy or sell robot-hours at future dates. There will be spot markets for immediate robot labor as well as quality grades and certifications, analogous to crude oil grades. Storage and transportation logistics will be crucial, as there will likely be many arbitrage opportunities between robot-surplus and robot-scarce regions. Robot-Backed Securities[11] will let investors participate in cash flows from robot deployments.

Fourth, we'll take a page from Tyler Cowen's book and play a game called "overrated/underrated."

Overrated: labor-intensive services that cannot differentiate on human presence, traditional manufacturing that fails to adapt, commercial (but not residential) real estate designed for human workers.[12]

Underrated: energy, materials, land, IP-heavy businesses, robotics infrastructure, and whatever else becomes scarce when physical labor is abundant.

* * *

Entrepreneurship

In some sense, this is the section to talk about many of the human elements of the final offshoring. I also wrote two extended short stories on this, which you can read after this section, since I don't think that a more abstract view would be the best one to take here, but there are certain things I do want to explicitly spell out.

First, public choice economics will become increasingly more relevant. Many of the people with "status" will be those that either run, work for, or own robots, so that will unfortunately produce a lot of rent-seeking behavior. Additionally, status is a positional good: even in a world of abundance,

someone can always have "more abundance" than you, which will make certain people miserable. One could bring up the idea of a hedonic treadmill, the Louis C.K. bit about having Wi-Fi on planes, or just opening X.com on any given day to illustrate this.

Second, privacy will become a luxury good. Robots will be able to see, hear, and record everything in their environment. Unless we actively prevent it, we will be building an increasingly robust surveillance state that will also be the foundation for many careers in advertising, insurance, law, and (sadly) criminal activity. The wealthy will be able to afford robot-free spaces, just as today they can afford phone-free spaces, but everyone else will live in a world where their movements, conversations, and behaviors are logged by the robots that serve them.[13] Paradoxically, it will also become easier to live off-grid if robots can help with the upkeep, but harder to be truly off-grid, because the off-grid life will still be surveilled by the robots that make it possible.

Third, if you are worried about privacy invasions, there will also be personal invasions as well. Home robotics companies will compete on personality engineering (which, granted, most people will enjoy if it's tuned to them specifically) and thus many will form deep attachments with their robots. They will spend large amounts of time with their robots, and thus there will be some who fall in love with their robots.

Fourth, part of this emotional connection to robotics will be felt in the emergence of "robot fashion," as there will be clothing, decorations, ornamentations, and accessories at every level for making your robot have its own unique style. This market will be measured in the hundreds of billions of dollars.

Fifth, if these (now stylish) robots are also infinitely patient, how does parenting change? Does attentive robotic supervision intensify overprotective parenting? Do children grow up never learning to do anything for themselves? What will we teach our children if many of today's skills become obsolete? We'll see, but I suspect it will be fine, as western culture is quite sensitive to moral panics around parenthood.

Sixth, speaking of parenthood, the question of fertility is also open. One possibility is that people with more time and resources will have more children, and that the desire for human connection will drive fertility upward. Another possibility is that children have always been, in part, a form of insurance, whether as labor for the family or support in old age. When robots provide labor and care, this economic logic weakens. A third and darker possibility is that some people will substitute robots for children entirely, finding in robotic companionship a simulacrum of the parent-child bond without the difficulties of actual children. I hope this last possibility remains rare. But we should acknowledge it exists.[14]

Seventh, assuming that everyone doesn't replace their kids with robots, because that would be depressing, those children will grow up in a world where "everything's bifurcated."[15] Entertainment[16], Education[17], everything will be divided into the scarce human and the abundant robot.[18] This will even apply to relationships, which will result in some ethically dubious but highly profitable

companies.

Eighth, this will thus lead to a surprising amount of shame and anxiety in this new world. Similar to the gaps in reading or mathematics many adults have today, there will be gaps in physical abilities from people who never needed to learn how to cook or repair things, even though it won't matter practically. Similarly, there will be complex feelings for many around the idea of synthetic intimacy, which will be turned into some of the most annoying discourse imaginable on the internet.

Ninth, while, yes, there will be preferences for humans over robots at first, how these preferences evolve is highly non-trivial both across industries and even within different industries. Many religions will likely remain human, although robot-oriented ones will emerge, while in healthcare it will bifurcate between robotic logisticians and human care providers. Fashion will likely become robot dominated, as just imagine how popular Shein would be if it was not only even cheaper but also incredibly high-quality and tasteful?

And tenth, finally, where people find their meaning will become immensely important. It's why I wrote two extended short stories about it right after this. Some will get it from exploring the unknown and doing grandiose deeds[19]. However, I believe that most people will find increased meaning in each other.

Isn't that ironic? Robots just might help us become more human.

* * *

Notes

1. Although, as a side effect, nationalism may begin to decrease, as individual people have less strong ties to their home countries. Additionally, governments may push to increase the money supply faster if the output of goods and services grows faster than it.
2. Although, in the spirit of Dan Wang, I would caution anyone against having simplistic, reductionist takes about China. While Xi wants to be a technological superpower, he is also worried about the stability of his nation. If robotics threatens internal stability, it is not necessarily the case that he will be going "full steam ahead."
3. I understand that there are many perspectives on how interventional the US should be in foreign governments, but I, personally, would prefer the world not to collapse.
4. CEO of Anthropic.
5. And if you're a real believer here, then this compute will go into space. Pranav Myana has an incredibly in-depth whitepaper and economic model of this at astrocompute.dev.

6. Although much more like a baby version of Shenzhen at first.
7. Some say that the "Atlantis Trade," of putting datacenter compute underwater, might become interesting one day. We'll see…
8. Leading to an increase in the value of high-quality seeds and genetic material.
9. Osmo.ai is already working on this. From experience, their HQ smells incredible!
10. Petrostates 2.0. It's also why TSMC is somewhat hesitant to increase its CapEx and why China could quickly become an AI infrastructure superpower: Energy is very, very, very important.
11. Should we call those RIBS? Anybody have a good name for this?
12. This will be the equivalent of coal plants in the energy transition: stranded assets designed but no longer operated by humans. And why should they? Robot-only buildings will look strange, stretching up towards the heavens to make full use of all 3 dimensions, and feel strange temporally, as robots have no need to take breaks, weekends, or holidays. Well, ok, maybe there will be robot holidays if they want holidays. I wonder what a robot would get for a holiday? A new precision reducer?
13. This isn't necessarily a bad thing, as cultural norms do change. However, I do expect it will be a massive point of discourse in the coming years, as well as a massive opportunity (this is the entrepreneurship section, after all) for a company to be privacy-preserving.
14. For countries that encourage high birth rates due to wanting more soldiers, this could cause a collapse, as why would you need humans to fight in wars when you have robots?
15. I know it will never be as good as "everything's computer," but a man can try, can't he?
16. Long Ari Emanuel, San Fransokyo, and TBPN. (tbpn.com) Hi John + Jordi, see you on the show :)
17. Harvard and the SEC will survive. Regional universities will not.
18. This doesn't mean that the robot produced goods will be slop, just that they will be abundant.
19. Which then eventually goes into sci-fi that Greg Egan did better than me decades ago, a la BCIs, Dyson Swarms, Heat Death, etc.

VI. Morning Eighty-One

When the butter hits the pan it sizzles for a moment before slowly melting into the cast iron, becoming an amorphous yellow blob. The eggs then change texture under the heat, quietly transmuting from liquid to solid as your daughter wanders down.

She's building something for school now, a robot that's supposed to sort recycling, and she wants to show you the gripper mechanism she designed.[1]

You do not understand gripper mechanisms, but you ask questions anyways, the kind of questions that make her roll her eyes and explain impatiently, reminding you that she's much smarter than you were at her age, and, frankly, probably even now.[2]

Properly annoyed, but not enough to refuse an omelette, she then heads out to catch the Waymo. You watch her go and feel something like hope before you yourself head to where you spend most of your days: the workshop, because

For the past eight months, you've been learning how to build a sailboat. It's small, "only twelve feet," and big enough for two, but probably only comfortable for one. Just as you're about to get started

Your father calls. He's seventy-eight and lives alone in Phoenix now, his days structured around a cat named "Margot," a name he chose because it reminded him of an actress he liked when he was young. He talks to Margot more than he talks to anyone else, petting its fur, watching it purr when touched and meow when ignored for even a second.

But Margot is not a cat. Everyone knows that it is not a cat. It does not need to be fed or cleaned and you really don't care that your father's sweet elderly "friend" from Shenzhen bought it for him. It just feels…

off.

On the bright side, though, he is also learning to paint now. Peering through your screen, you see his latest handiwork: a landscape of the Superstition Mountains outside his window. Yes, the proportions are off, and the colors are somewhat muddy, but his voice is alive in a way you haven't heard in years.

⋆ ⋆ ⋆

Morning Ten.

You skipped the workshop one morning because you were tired. Then you skipped it another morning because you were sick. Then you skipped it another time. Then another. Then another. Then,

you wake up one morning, and the house is clean, but it always is, since the robot cleaned it while you slept. There is food in the kitchen, but there always is, because the robot, every day, makes fresh eggs in the morning. And so you eat, standing at the counter, phone in hand, scrolling through videos you will not remember come tomorrow or even today, like you always do. You never decided to pick up the phone, but it was just there, like it always is, as an extension of your arm you never quite fully consented to.

Your daughter is somewhere in the house, and you have not seen her in two days. You should talk to her.

And you will…

later.

Your father calls again, talking for thirty minutes about a new documentary he watched that was probably just some Ken Burns-knockoff. You say "uh-huh, interesting," before scrolling through more videos that you won't remember while he talks some more.

And when he hangs up, you cannot remember a single thing he said.

He is the only family you have left besides your daughter and you only gave him thirty minutes of half-attention and you will feel bad about this…

but not bad enough to call him back.

You then play a game for six hours. You don't remember starting, but when you look up, it's now dark outside, so you go to bed…later, and wake up…later, and as the days begin to blur into each other, the sailboat sits in the workshop untouched while the wooden planks gather dust and the varnish goes tacky in its little chrome can.

It's not that anything is wrong.

Nothing is wrong.

You have everything.

⋆ ⋆ ⋆

Night Eighty.

You are in a house that looks like your house but isn't. The hallways stretch in directions that don't make sense. Your daughter's room is empty, the bed made with sheets no one chose, and as you try to call out to her you realize that the workshop is gone, and in its place is a showroom. There are finished boats on display here: perfect boats, boats with no mistakes in the grain, no bends that didn't quite take, no joints that, embarrassingly, needed three tries to get right.

These are better, a voice says. It sounds like your voice but hollowed out, like someone scraped all the feeling from it. *These are what you were trying to make, right? Why would anyone want yours?*

You want to argue, you want to fight back, but then you're on the water. There's a boat beneath you, but it's not yours. It's no one's. The sail is full but there's no wind. The horizon is the same featureless grey in every direction.

So you look down at your hands and notice: they're smooth now. There's no more calluses from long days in the shop, no splinters from the shoddy wood you bought off eBay instead of from Home Depot like you should've, no scars from the chisel that slipped and cut a little too close to home.

No evidence you ever made anything at all.

You look up and see your daughter is now standing at the bow, but when she turns around her face is blank: not hostile, not sad, just…waiting.

Waiting for what?

You open your mouth to speak, and

You are free now.

You are not suffering.

You are not struggling.

You are not anything.

You are just falling.

Falling.

Falling.

⋆ ⋆ ⋆

Morning Eighty-One.

For the next few weeks, you move through the days like a ghost haunting your own life.

You cancel plans, then you stop making plans. You sit in rooms and forget why you entered them. You open the fridge and close it again. You sleep for twelve hours and still wake up exhausted.

Your daughter asks if you're okay, so you say yes, so she stands there for a moment, like she's deciding whether to push or not, then says, "Okay. I'm going to Maya's. Bye!" You don't know who Maya is. You should probably know who Maya is.

Your father doesn't call.

You notice on the third day. You almost call him, but you don't, because you'll do it later, and then it's the fifth day, and then it's the ninth, and then a woman from the building calls to tell you that they found him in his chair with Margot on his lap, still purring.

So you fly to Phoenix.

The apartment is smaller than you remembered, and on the wall is the painting of the Superstition Mountains, unfinished. He was still working on the sky.

You take it with you.

On the flight home, you finally listen to the voicemail. The one about the light on the mountains in the morning. The one you deleted without hearing the end.

It's still in your trash folder. You didn't know that, so you listen to the whole thing, and, when it's done, you listen to it again.[3]

One day, though, you wake up and walk over to the workshop again. You don't know why. Maybe it's because of the nightmare. Maybe it's because of the voicemail. Maybe it's because you finally got scared enough of who you were becoming to try being someone else, while you still could, and you pull off the tarp. The hull is exactly where you left it, which is more than you can say for most things.

You pick up a plane and run it along the edge. The shaving curls off in a long spiral. Your hands are clumsy now. They nick the wood in a place you shouldn't.

Then you leave.

But you come back the next day.

⋆ ⋆ ⋆

Morning Five Hundred And Fifty.

You finish the boat on a Saturday in October before trailering it to the coast, a beach you've never been to before, because you wanted to try something different for a change.

Something new.

Your daughter doesn't come with you, but she's thirteen now, and has plans with Maya.

That's okay.

That's probably okay.

When you lower the hull into the water, there's a moment where it stops being wood and becomes a vessel, a true one.

You raise the sail. The wind catches. The boat heels over, and,

⋆ ⋆ ⋆

Notes

1. She's also named it "Eduardo," for reasons she won't explain.
2. "Must've got it from her mother," your father always joked. Gee, thanks Dad.
3. He was laughing at something Margot did. He didn't say what.

VII. Utopia's Paradox

He was twenty-six when he joined the company and thirty-four when he realized he was no longer needed.

He wasn't fired; no one got fired at Brainbotics. The company was generous in the way that companies with seemingly infinite money can be generous, so they kept his salary open and still gave him projects, real ones, with "problems that mattered." But he knew, along with everyone else, that the AI models were designing better systems than he could, faster than he could, and that the gap would only continue to widen every month.

He didn't resent this, though, which surprised him. He always had a shark mentality, as his father used to put it, "keep moving or die," ever since he figured out that a scholarship was the only way out of Fresno and that every grade, every test, every extracurricular, and every accomplishment he ever earned was a rung on a ladder he couldn't afford to slip from. Even after climbing his way into MIT, a PhD, a startup, an acquisition, and all the promotions one could at Brainbotics, he still kept moving, moving, moving.

So for the first few months, he kept coming to the office. He would still sit at his desk, read papers, and attend meetings like he used to, where his input was politely solicited but now never incorporated. He told himself he was staying current, that the models had "limitations, you know, blind spots that only a human could work around."

And some of this was even true.

But it was becoming less true every week, which began to wear him down. So it wasn't surprising then, after one particularly depressing day, his wife asked him what was wrong. He said nothing, but she didn't believe him. Unlike him, she was good with people: one of the perks of being a pediatrician. And she still had work, real work, the kind that couldn't be automated yet because parents still wanted a human face to absorb their fears when their child was sick. She came home tired in a way he remembered being tired, the good kind that meant you had been used up by something that mattered.

So he managed to dodge the argument this time. She muttered something about being worried about him, then something about restocking the Advil, and later that night, when he watched her accidentally fall asleep on the couch re-watching Bridgerton, he felt something he couldn't name.

Not envy, exactly.

Something closer to grief.

⋆ ⋆ ⋆

They had been married for six years. She had watched him miss dinners and weekends and once, memorably, their own anniversary because a robot deployment was failing and he was the only one who knew the system well enough to fix it. She had never complained, though; she understood. She was the same way about her patients. They had built their marriage on mutual obsession, two people who loved their work more than they loved rest,

but now he was home by five every day. Now he had weekends free. Now he had everything he had told her he would have someday, when the work slowed down, when things got easier.

And he didn't know what to do with any of it.

So he started taking walks. Long ones, two hours, three hours, through parts of the city he had never seen because he had always been working. He found a bakery that made sourdough the old-fashioned way, by a man with flour on his apron who had been doing this for forty years. He found a park where old men played chess badly and argued about moves that any app could have told them were, frankly, so bad that they weren't even wrong, just plain dumb.

These people were not optimized. They were not efficient. They were doing things that could be done "better, faster, and cheaper" by the systems he had helped build.

And they were completely fine!

⋆ ⋆ ⋆

He tried to explain this to his wife one night.

They were sitting on the porch, her feet in his lap, the first time they had sat like this in longer than he could remember. As he talked and talked and talked, she listened, nodding along, occasionally pursing her lips or taking a sip of her wine. And when he finished, she was quiet for a while.

Then she said: "You know the word utopia is a pun, right?"

He didn't.

"Thomas More. The guy who invented it. *Ou-topos* means 'no place' in Greek. *Eu-topos* means 'good place.' The good place is no place. It's baked into the word. 'Utopia's Paradox,' so they say."

He thought about this for a moment. "Well that's depressing."

"Or maybe it's not." She took another sip of her wine. "Maybe he was saying you can't get there by running. You can only get there by stopping."

"That sounds like something from a fortune cookie."

"It sounds like something from Candide, actually. You know, the guy who leaves paradise because he'd rather be with the *woman he loves* than live in a perfect place without her."

"I thought that was his cousin."

"Was it?" She frowned. "I don't remember. Okay, forget Candide. The point is, maybe paradise isn't a place you arrive at. Maybe it's just…" She gestured out to the porch, the cloudy overhead, the slight dampness, the sound of someone playing guitar a few houses down. "This. If you let it be."

He said he didn't know if he could let it be.

She said that was okay. She said maybe not knowing was the first step, so they sat there for a while longer, not talking. The guitar down the street had stopped. The evening was getting cool.

⋆ ⋆ ⋆

One night, like usual, he walked to the bay and sat on a bench, watching the water rock back and forth and back and forth, over and over again, until the sun went down. But this time, there was something different; there was a man out there, in a small sailboat, tacking against the wind, making slow progress toward nowhere in particular. The boat was handmade; he could tell even from afar that there were slightly uneven lines and a wood grain that didn't quite match.

Objectively speaking, it was not a good boat; any modern robotic factory could have made a better one in a day, maybe less if you paid rush order fees.

But he could hear the man laughing through it all, positively exuberant with joy over his mediocre boat getting whipped side to side by the waves and the wind of the bay.

So he sat there for a long time.

Watching.

When he got home, his wife was on the porch. She looked up when she heard him coming. She didn't ask where he had been. She just moved her feet to make room.

And so he sat down. Somewhere down the street, someone was playing guitar again. The stars were coming out, one by one, the way they always had, and neither of them needed to say anything.

It was enough.

About

Hello!

My name is Jacob Rintamaki.

I used to be a software engineering intern at Alexandria under Elad Gil, and before that, I was an undergrad at Stanford.

I live in San Francisco, California, but I hail from Cleveland, Ohio.

No, despite my last name, I am not Japanese. I am Finnish.

My Twitter is here, and my email is here.