


The Industrial Explosion

Tom Davidson & Rose Hadshar

May 2025



Contents

The Industrial Explosion

| | |
|--|----|
| The industrial explosion will start after the intelligence explosion, and will proceed more slowly | 4 |
| Three stages of industrial explosion | 5 |
| AI-directed human labour | 6 |
| Fully autonomous robot factories | 7 |
| Nanotechnology | 8 |
| How fast could an industrial explosion be? | 8 |
| Initial speed | 9 |
| Acceleration | 10 |
| Maximum speed | 12 |
| Appendices | 13 |
| How fast could robot doubling times be initially? | 13 |
| How quick might robot doublings become by the time we reach the earth's carrying capacity? | 16 |

Tom Davidson & Rose Hadshar

Tom did the original thinking; Rose did later thinking and writing.

For correspondence, please email contact@forethought.org.

This is a rough research note – we’re sharing it for feedback and to spark discussion. We’re less confident in its methods and conclusions.

Once AI can automate human labour, positive feedback loops could lead to a rapid increase in AI capabilities. So far, most analysis has focused on increases in *cognitive* capabilities - an ‘[intelligence explosion](#)’. But *physical* capabilities could also increase rapidly, in what we’re going to call an ‘industrial explosion’. Sufficiently advanced robotics could create a feedback loop where robot factories build more and better robot factories which build more and better robot factories, and so on.


The incentives to push towards an industrial explosion will be huge. Cheap abundant physical labour would make it possible to alleviate hunger and disease. It would allow all humans to live in the material comfort that only the very wealthiest can currently achieve. And it would enable powerful new technologies, including military technologies, which rival states will compete to develop.

The speed of the industrial explosion matters for a few reasons:

- Some types of technological progress might not accelerate until after the industrial explosion have begun, because they are bottlenecked by physical infrastructure for running experiments
- If the industrial explosion is fast, a nation that leads on AI and robotics could quickly develop a decisive economic and military advantage over the rest of the world.
- All things equal, a faster industrial explosion means the AI could overthrow humanity sooner – many routes to AI takeover go via AI controlling new physical infrastructure.
- The industrial explosion will feed back into the intelligence explosion, by rapidly increasing the world’s supply of energy and compute.

This post presents an initial analysis of the dynamics of the industrial explosion. We argue that:

- The industrial explosion will start [after the intelligence explosion starts](#), and it will initially proceed more slowly than the intelligence explosion.
- Schematically, we can think of the industrial explosion unfolding in [three stages](#) :
 - **AI-directed human labour**, where AI-directed human labourers drive productivity gains in physical capabilities.
 - **Fully autonomous robot factories**, where AI-directed robots (and other physical actuators) replace human physical labour.
 - **Nanotechnology**, where physical actuators on a very small scale build arbitrary structures within physical limits.
- The industrial explosion could ultimately become [extremely fast](#), with the amount of physical labour potentially doubling in days or weeks.



Three stages of the industrial explosion.

The industrial explosion will start after the intelligence explosion, and will proceed more slowly

The industrial explosion will likely start after the intelligence explosion because physical tasks will be automated after cognitive tasks. Cognitive tasks are easier to automate for a few reasons:

- Cognitive tasks, especially for key domains like AI R&D, are less wide-ranging.
- There's much more data on how tasks are completed (like documents, meeting notes, potentially computer screen recordings). By contrast, it's expensive to gather data for tasks in the physical world.
- The tasks are entirely virtual, avoiding many tricky real-world frictions.

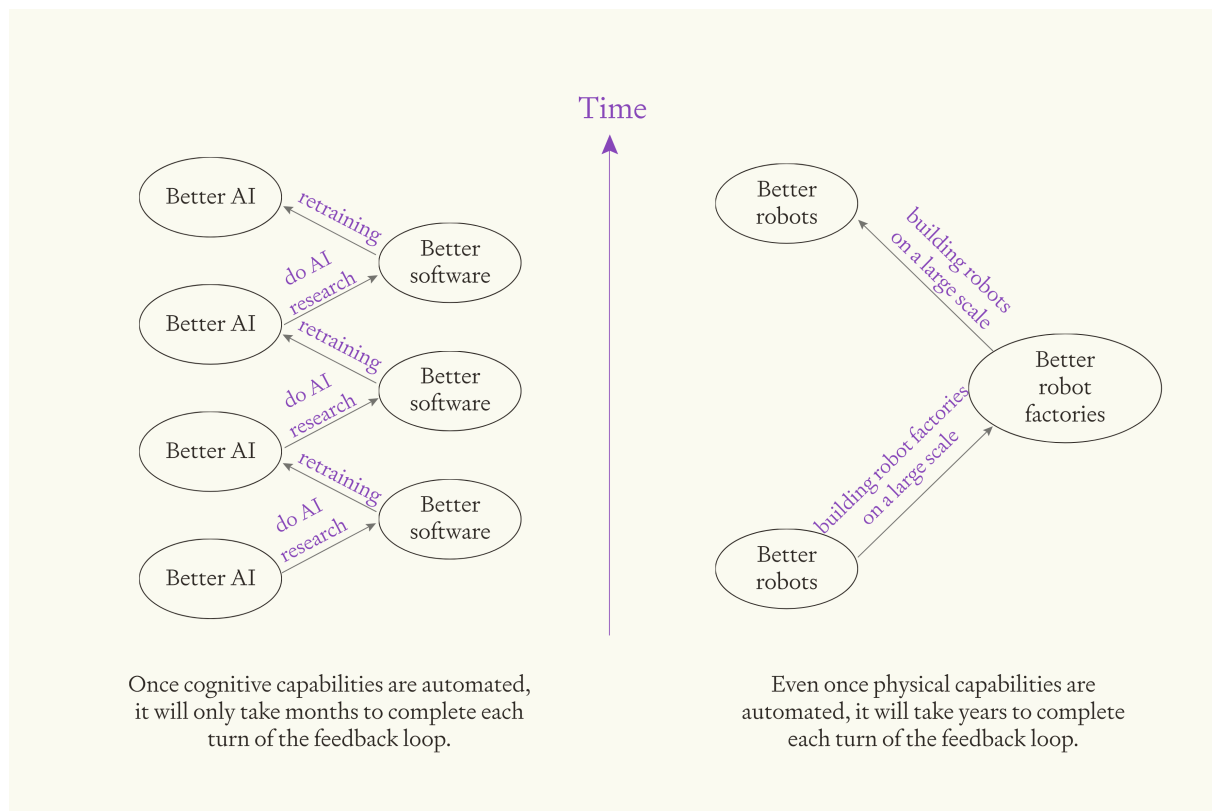
As well as starting later, the industrial explosion will also be slower than the intelligence explosion.

The first reason is that the *current* rate of technological improvement for AI cognition is faster than the rate of technological improvement in robotics. AI chips double in FLOP/\$ every ~2 years. AI algorithms double in efficiency every year or less. We think that robot technology

doubles in efficiency more slowly than this, perhaps every 1-4 years.¹ So the technologies that will drive the intelligence explosion are increasing much faster than those that will drive the industrial explosion.

The second reason the industrial explosion will be slower is because the feedback loop of “robots make more robots” is has a bigger time lag than the feedback loop for “AI makes smarter AI”:

- As soon as we train a powerful AI system, we’ll be able to run millions of copies in parallel. By contrast, once we develop perfect humanoid robots it will take years before we produce millions.²
- It takes longer to build a complex factory from scratch than to train an AI system from scratch.³



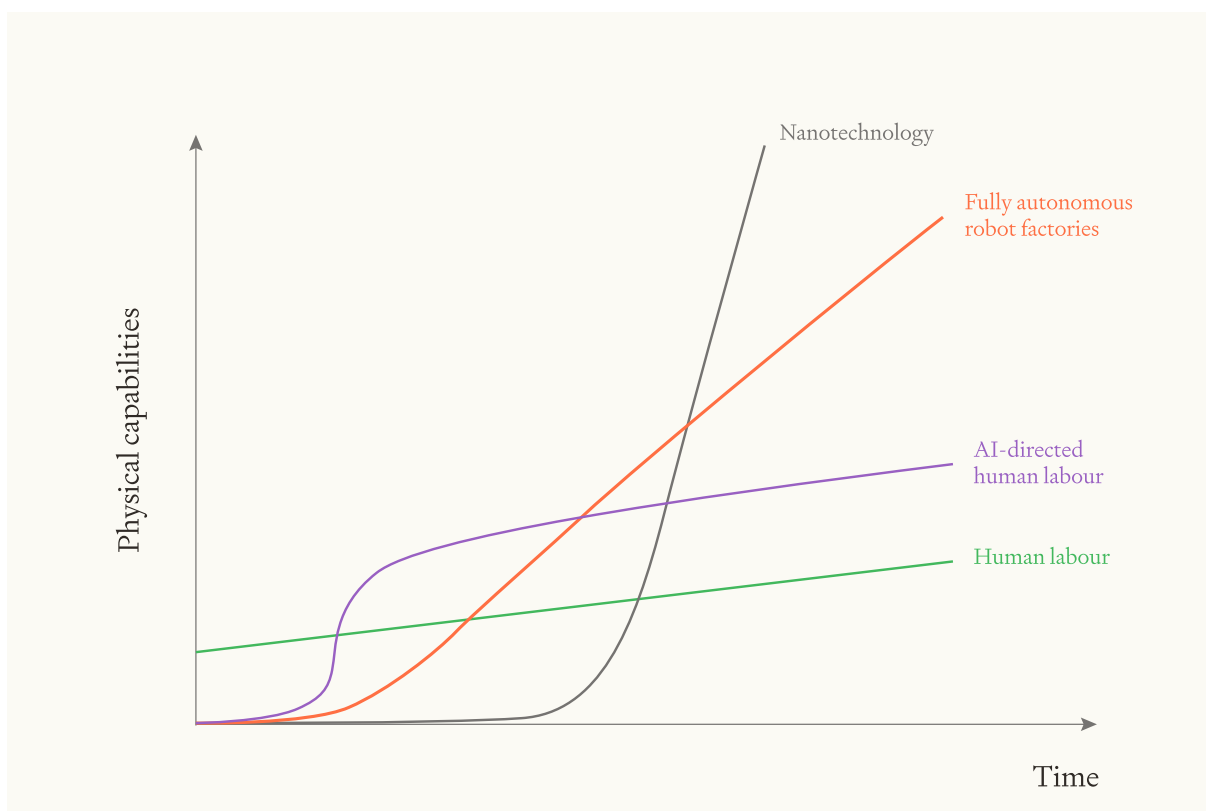
So the industrial explosion will start after the intelligence explosion, and happen more slowly.

Three stages of industrial explosion

-
- 1 Goldman Sachs [estimates](#) that robot costs have recently fallen by 40% per year (1.4 year efficiency doublings), but had previously forecast cost reductions of only 15-20% per year (3-4 year efficiency doublings).
 - 2 See [Todd \(2025\)](#) for estimates. Returns from AI-directed human labour would come online more quickly than this.
 - 3 Training an AI system currently takes [months](#).

Schematically, we can think of the industrial explosion unfolding in three phases:⁴

- **AI-directed human labour**, where AI-directed human labourers drive productivity gains in physical capabilities.
- **Fully autonomous robot factories**, where AI-directed robots (and other physical actuators) replace human labour.
- **Nanotechnology**, where physical actuators on an atomic scale build arbitrary structures within physical limits.



Three stages of the industrial explosion.

AI-directed human labour

In the first phase, AI-directed human labour will drive large gains in the productivity of physical production.

Today, human physical labour is not maximally productive:

- Some workers are much more productive than others. You can see this from differences in salary worldwide and within countries, which are significant both between individuals and across one person's career.
- Even the most productive workers might not be at the physical limits of human productivity.

⁴ None of these phases are inevitable: humans may choose not to pursue explosive industrial expansion beyond a certain point. But absent human intervention, we expect an industrial explosion to progress through each of these stages.

AI could bring the economic productivity of human manual workers close to or beyond the productivity of the very best human workers today.⁵ For example:

- Human manual workers could wear sensors (e.g. phones with cameras and microphones) which allow real-time AI monitoring of their actions. With this vast amount of data, AI could generate specific real-time advice to each worker.
- Using this data, AI could generate process improvements which increase the efficiency of whole factories and industries.
- AI could also coordinate the actions of many disparate humans on complicated projects by tracking all their actions in real time and adjusting plans accordingly.

Because AI-directed human labour only requires advances in cognitive capabilities, this phase will probably happen before fully autonomous robot factories or nanotechnology. It could in principle be rolled out quite quickly, though in practice this will depend on human adaptability, regulation and other human factors.

This phase will involve lots of humans doing physical labour, as their cognitive labour is no longer useful.⁶

Fully autonomous robot factories

After increasing the size of the physical economy by a moderate factor,⁷ AI-directed human labour will run into natural limits: humans can only work so efficiently.

At that point, further demand for physical labour could drive the development of robots and other physical actuators that can fully automate human physical labour.

In practice, physical labour will become increasingly automated in a gradual way:

- Today, physical labour is done by a mixture of humans and specialised robots, and the process of production is ultimately directed by humans.
- In the phase of AI-directed human labour, the process of production will come to be directed by AI systems.
- The proportion of physical labour done by humans will begin to fall, as physical capabilities increase and better robots (and other physical actuators) are produced.
- As robots continue to improve, humans will only perform tasks where humans have strong comparative advantage.
- Eventually robots will be better than the best human manual labourers at all tasks.⁸

5 We're not claiming this would necessarily be good - increasing human productivity might come at the cost of freedom and other things we hold dear. We're just claiming that this will be possible once AI gets good enough.

6 It's possible that wages for physical labour get higher at this point, though if the supply of physical labour increases a lot this is less clear. The human wage distribution might also get flatter, if AI is directing human physical motions such that humans become more substitutable.

7 See [below](#) for thoughts on how large this factor might be.

8 This seems technologically feasible. Humans are proof-of-concept that human-level physical capabilities are possible. It would be surprising if evolution, which is a blind search process, had reached the physical limits of physical capabilities. Besides, variation among humans, and animals and specialised robots which outperform humans in specific areas, show that improvements on human capabilities are possible.

Of course, humans may choose not to fully automate physical labour. But absent human bottlenecks, economic incentives and increasing physical capabilities would eventually lead to robots (and other physical actuators) that can fully replace human workers.

If physical labour is fully automated, then an array of AI-directed robots and other physical actuators will be able to autonomously do all economic tasks, including making more robots. In other words, the robots can self-replicate. This is important, as it creates the positive feedback loop that's required for an industrial explosion.

Indeed, autonomous robots may initially be specialized for the purpose of making more robots over other tasks, because this task will be so economically valuable.

Sometimes 'self-replicating robots' is used as a shorthand for these AI-directed physical actuators. But it's important to realise that:

- Individual robots probably won't be able to replicate themselves. It's more plausible that there will be a whole array of actuators in a set of factories, each producing parts for machines which produce parts for machines... which are collectively capable of self-replicating.
- The physical actuators won't just be humanoid robots. Initially, many of the actuators may be humanoid robots, as most physical equipment is currently designed for human labourers.⁹ But ultimately there will likely be large efficiency gains from relaxing the constraint of human compatibility, and producing physical equipment and robots which are more optimised.

Nanotechnology

Eventually, fully automated physical labour will run into physical limits: it won't be possible to build physical objects any faster.

But smaller objects are faster to build. We see this empirically, with bacteria and other small organisms self-replicating faster than larger organisms. There are also basic engineering principles which support this conclusion (for example, smaller objects have a bigger surface area to volume ratio, so can absorb more materials per unit mass).¹⁰

Because smaller objects are faster to build, there will be returns to designing smaller and smaller machines, with faster and faster throughput.

In the limit, an industrial explosion could enter into the third phase, nanotechnology, where physical actuators on a very small scale build a very wide range of structures.¹¹

How fast could an industrial explosion be?

The speed of the industrial explosion will likely change over time. We can consider:

- How fast will the industrial explosion be *initially*?

9 Though human workers may still operate those machines while new machines are made for the robots.

10 See Drexler, *Nanosystems* (1992) for an extended analysis.

11 For an overview of how nanotechnology might arise and might transform the physical world, see Drexler, *Radical Abundance* (2013).

- How quickly will the industrial explosion *accelerate* ?
- How fast will the industrial explosion *eventually become* ?

It's hard to make substantive claims about the speed of the industrial explosion, as it requires making so many assumptions. Nevertheless, we can make some general claims.

Initial speed

One-time gain from AI-directed human labour

For the first phase, AI-directed human labour, we could operationalise the speed of the industrial explosion in terms of productivity.

How large an increase in total productivity might AI-directed human labour give?

AI direction could make most workers much closer in productivity to the best workers. The difference between the productivity of the average and the best manual workers is perhaps around 2-6X:

- In the US, labourers earn [\\$13-26](#) an hour.
- Studies of the variation in output *within a given firm* suggest the very best workers are 1.5-2X more productive than the mean.¹²
- AI will also make whole firms more productive by improving planning and organisation. Top firms are ~3X more productive than the mean.¹³

We should round this up further though, to account for the possibilities that:

- AI leads to performance improvements beyond that of the best human workers today.
- There is a one-time increase in total productivity from large numbers of cognitive labourers switching to manual labour. This could easily be another 2X.

We're uncertain about how large these uplifts might be, but it looks like – combining the gains from more productive individual workers, more productive firms, and more total human workers – the overall increase in physical output here might be about 10X.

| Factor | Increase in physical output |
|-------------------------|-----------------------------|
| More productive workers | ~2X |

12 [This](#) study finds that for a selection of mostly manual jobs, one standard deviation from the mean in performance corresponds to 1.2X of mean performance. Assuming normality, the 99.97th percentile of workers (three standard deviations from the mean) would be 1.6X mean performance. [Another similar study](#) finds three standard deviations would correspond to a 1.9X increase.

13 The standard deviation in firm productivity from [these three papers](#) suggests that improving firm productivity by two standard deviations would increase productivity by 1.7X - 3.7X. (We consider two standard deviations rather than three in this case because, while AI will improve many determinants of firm productivity quickly, without advanced robotics it will have a much smaller and slower impact on other determinants like the quality of physical equipment.)

| | |
|--------------------------|------|
| Better run organisations | -3X |
| More manual workers | -2X |
| Total | -10X |

Initial doubling times for autonomous robots

As the industrial explosion transitions from AI-directed human labour to increasing and eventually full automation of physical labour, we can start to operationalise the speed of the industrial explosion in terms of robot doubling times: the time it takes to double the number of robots (and other types of physical actuators) in the world.¹⁴

The most recent doubling in the number of robots in the world took [6 years](#). It's hard to say how quickly self-replicating robots could double in number, but in [an appendix](#) we use a couple of approaches to tentatively estimate that with current physical technology (but abundant AI cognitive labour) this might be on the order of a year, rather than a month or a decade. It could be faster still if AI can quickly drive rapid technological progress without an industrial explosion happening first (for example, by quickly developing advanced nanotechnology).

Acceleration

If robot technology remains constant, the growth rate of robots and other physical actuators will be constant (ignoring resource constraints for simplicity)¹⁵.

But if technological improvements mean that robots become *twice* as easy to make, then the growth rate will double.

Ideally, we'd get data on how much you need to increase the stock of robots and other physical actuators before their price halves – an [experience curve](#) for robots. We don't have trustworthy data on that unfortunately. But there are many papers estimating this quantity for related sectors:

14 During the transition to full automation of physical labour, robot doubling times will not initially map to the speed of overall industrial expansion, as there will likely be far fewer robots than humans. But at some point (probably prior to full automation), there will be a large enough number of robots that robot doubling times substantially drive the speed of industrial expansion.

15 It's possible that resource constraints are a bottleneck, and this is an important area for further research, but our guess is that they won't be. Historically, resource bottlenecks have never capped GDP growth – they've been circumvented through a combination of efficiency improvements, resource substitutions, and improved mining capabilities. There is no reason to expect resource bottlenecks to suddenly bite just as robotic physical capabilities surpass those of humans. Indeed, economists are well aware that claims that growth will be bottlenecked by some essential resource have often been [wrong historically](#) and should be treated with suspicion.

In the technological limit, the human body shows that it's technologically possible to make "humanoid robots" with abundant materials like Carbon, Oxygen, Nitrogen etc, such that the limiting factor on quantity will be energy not resource bottlenecks. (And there's room for a [3-4 OOM increase in world energy use](#)). That doesn't preclude resources being a bottleneck along the way, but is some additional evidence that they won't be.

| Sector | How many times production must double to halve the cost | Source |
|--------------------------------------|---|--|
| Chemical industries | 1 - 10 | Nagy et al (2013) , Supporting Information 1 |
| Hardware industries | 1 - 2.5 | Nagy et al (2013) , Supporting Information 1 |
| Energy industries | 2 - 10 | Nagy et al (2013) , Supporting Information 1 |
| Other industries (mostly electrical) | 2 - 5 | Nagy et al (2013) , Supporting Information 1 |
| Aggregate economy | 3 | Bloom et al (2020) , Table 7 ¹⁶ |
| Moore's law | 0.2 | Bloom et al (2020) , Table 7 ¹⁷ |
| Agricultural sectors | 2 - 10 | Bloom et al (2020) , Table 7 ¹⁸ |
| Robots | 1 | ARK Invest (don't provide raw data) |

So if robot technology improves with the same learning curve as the aggregate economy, it will take 3 doublings before the cost of robots and other physical actuators halves and (as a consequence) the robot growth rate doubles. If it's like Moore's law, then it will accelerate much more quickly, but that is famously an outlier.

Overall, it looks likely that the number of robots will double 1-5 times before the robot growth rate doubles.

16

17

18 This paper measures production not as units produced but as cumulative R&D spending. In the context of robot doublings, that corresponds to the assumption that each time our robot population doubles, we'll also double the total amount of robot R&D. That seems roughly reasonable, as we'll have twice as much opportunity to gather data from those robots and run experiments with them.

Maximum speed

We can upper bound how fast the industrial explosion could become by thinking about how fast robot doublings could become in the limit of technological feasibility (though human bottlenecks might cause us to move more slowly).

To sustain physical growth that would be valuable to humans, the self-replicating machines need to be complex enough to make the machines that make the machines... that make all machines in the modern economy. They might bootstrap by having biological instincts that, under certain conditions, cause them to stop replicating and instead start making increasingly complex machines. Alternatively, they might be configured so that they can receive instructions from AIs who would then direct their behaviour so that they build the desired machines.

How quickly could such machines replicate? One way to estimate this is to look at biological analogies.

Some [bacteria](#) can double in hours. But these organisms are very simple and [cognitively basic](#) so may be unable to bootstrap to complex machines.

Instead, we can look at the fastest doubling times for biological organisms which have brains, and therefore may be capable of executing sophisticated behaviour based on their sensory inputs. In optimal conditions, fruit fly populations can double in days.¹⁹ This is proof of concept that biological replicators with brains can double in days.²⁰

Still, fruit flies are physically weak and cognitively fairly basic; perhaps they are too limited to rebuild the full physical economy. Rats are a more conservative example, and in good conditions they can double in about 6 weeks.²¹

One source of scepticism here is that the earth can only carry so many robots, and we might reach the limit before robot technology becomes good enough for such quick doublings. But a [quick BOTECH](#) suggests that, extrapolating the experience curves discussed above, we would get doubling times of less than a day before reaching the earth's robot carrying capacity.

It seems reasonable to use days or weeks as an upper bound on how fast robot doublings could become, based on biological analogies. This is very fast indeed.²²

Thanks to Owen Cotton-Barratt, Max Dalton, Oscar Delaney and Fin Moorhouse for helpful feedback.

19 At 29°C, it takes [7 day](#)s from egg to adult, and total lifespan is around [40 days](#). Females can lay up to [100 eggs](#) per day, from the day they become adults. Naively, a population of 2 fruit flies (one male and one female) would become a population of 202 adults in 8 days (~1 day doubling time).

20 Still, fruit flies may still not be able to follow arbitrary instructions from AIs. Computer chips can execute any software that is written in machine-readable code. This allows anyone to write software to their pleasing, translate it to machine-readable code, and then use the computer chips to run the software. This isn't possible with fruit flies because there is no "fruit-fly-readable code" – no flexible way to describe software that any fruit flies' brain can interpret.

21 [NCERT Class XII Biology, Chapter 11 "Organisms and Populations," p.194](#) .

22 Eventually, increasing the number of robots will be constrained by access to matter and energy, and will need to slow until we can expand beyond earth to gather more of these resources. The closer we are to carrying capacity when an industrial explosion begins, the less time there is to accelerate, and the lower the maximum speed will be.

Appendices

How fast could robot doubling times be initially?

Once AI can manipulate robots as well as humans can manipulate their bodies, how fast will robot doubling times be?

Here, we give a preliminary sketch of a rough order of magnitude estimate. We assume physical technology is the same as it is today, but assume that there is cheap and abundant AI cognitive labour to control robots and other types of physical capital.

We use two separate estimation approaches, though both have significant uncertainties.

| Estimate | How it works | Bottom line |
|--|---|-------------|
| How fast is physical capital at making more physical capital? | Look at a factory that makes more factories. Doubling time = (value-added each year by the factory) / (value of the physical capital in the factory) | -1 year |
| How long would it take a humanoid robot to pay for its own construction? | Doubling time = (wage of a productive manual worker) / (cost of making a humanoid robot) | -1 year |

How fast is today’s physical capital at making more physical capital?

(Thanks to Constantin Arnscheidt and Damon Binder for raising this approach to our attention.)

Self-replicating robots will involve a wide variety of physical capital – e.g. factories, machines and infrastructure – making more physical capital. So one question is, how quickly can today’s physical capital produce more physical capital?

We can estimate this by comparing the \$ value of the physical capital in a particular factory to the value that factory produces in a year (in the form of new physical capital).²³ For example, if a \$1b factory produces \$1b of value each year, then that suggests the total amount of physical capital stock could double in a year. If it only produces \$0.5b of value, then a doubling would take 2

23 In economics, this ratio is referred to as value added / PP&E – the \$ value added of a factory divided by the cost of the Property, Plant, and Equipment.

years.²⁴

According to data from the Bureau of Economic Analysis the US manufacturing sector produced [\\$2.6bn](#) of value in 2022 using [\\$5.4bn](#) of physical capital.

These numbers naively suggest that self-replicating robots could double in about two years.

This estimate is a bit aggressive for a couple of reasons:

- Complex factories often take more than a year to build. Building them faster might cost more. This is a small factor as very ambitious factories can be [built in less than a year](#) when people are really trying.
- It ignores the human labour costs of producing physical capital. These are very roughly half of the costs, though robot costs will likely be lower. We think this factor will be outweighed by the fact that fully autonomous factories can operate day-and-night, unlike normal factories.

On the other hand, the estimate is very conservative in ignoring productivity improvements from abundant AI cognitive labour. Having physical capital (and robot labour) be controlled by superhumanly smart and motivated AIs could significantly boost productivity. This might reduce the doubling time, by a factor of 2-4X.²⁵

So, all in all, this first approach suggests an initial robot doubling time of roughly 1 year or less.

How long would a humanoid robot take to pay for its own construction?

To begin with, we'll think through an unrealistic but simple hypothetical scenario. Then we'll consider how this might transfer to the real world.

Imagine a hypothetical where self-replicating humanoid robots drop from the sky tomorrow. They can perform all physical tasks as well as a human, cost the same as today's robots, and are manipulated by a limitless supply of AI systems who direct them as well as the most productive human workers. These robots go on to self-replicate without any help from humans.

A very basic economic analysis suggests a robot doubling time of ~5 months:

- Human manual labourers in the US earn [~\\$40k](#) a year. But these robots will be 6X more productive as they work day and night (3X) and are more productive than average ([2X](#) ²⁶). So that's ~\$240k/year.
- Humanoid robots currently cost [~\\$30-150k](#) to produce commercially, so let's say ~\$100k.
- This suggests that the robots could pay for their own construction every ~5 months.

24 Growth in physical capital is much slower today because a small fraction of physical capital is used to make more physical capital and because factories require human operators, which would bottleneck growth.

25 Above, in the section "One time gain from AI-directed human labour", we estimated a 2X gain from having individual human workers become more productive and a 3X gain from firms being run better. But a 6X reduction in doubling time seems aggressive given the difficulty of quickly making new factories.

26 This study finds that one standard deviation of performance in dollars corresponds to being 2x more productive than the median. Note that they find the second standard deviation is 7x more productive, and the third standard deviation is 25x more productive, so this factor could be higher.

This basic analysis suggests that in our hypothetical, robot doubling times will be on the order of months.

But this is too simplistic. There are two strong reasons to expect that the doubling time (even in our unrealistic hypothetical) would actually be longer:

- **It's more expensive to build factories from scratch than to rent them.** The cost of making a robot (~\$100k) already takes into account the costs of renting the factory space. But at scale it would be necessary to build new factories,²⁷ and this is much more expensive. Mathematically, we can think about this as a multiplier on the amortized construction cost of the physical capital used to build the robots. The multiplier equals the time over which the physical capital construction is currently amortized, divided by the robot doubling time. If the cost of building factories is currently amortized over 10 years, then a 12 month doubling time would increase these costs by 10x.
 - A 2 year doubling time seems compatible with this factor, as robots could produce 6X as much value as in 5 months, and the amortised costs would only be 5X higher.
- **It is more expensive to make things quickly than to make them slowly.** Currently it takes [1-2 years](#) to build a new factory. But to double the robot population in 5 months, new factories would need to be built in weeks, which is a 50X speedup. We should expect very large cost penalties for this, and it probably requires significant technological progress for this to even be possible.
 - Again, the penalties for a 2 year doubling seem like they wouldn't be large.

These factors stop biting at around 2 years, so shift our estimated doubling time up to 1 - 2 years.

The hypothetical doubling time might be shorter due to **lower labour costs**. Currently, robot construction costs include paying for human cognitive labour. In our hypothetical scenario, there is abundant AI cognitive labour, so these costs don't need to be paid. If half the cost of robot construction is currently human cognitive labour, this would be a 2x reduction in the doubling time.

This leaves our estimated doubling at about 1 year or less in our hypothetical scenario. We put less weight on this method than the one above.

How does this translate to the real world?

There are a few related reasons to think that once physical labour is fully automated in the real world, initial robot doubling times might be towards the shorter end of that range:

- In our hypothetical, humanoid robots self-replicate without any help from humans. In the real world, humans will be helping to make more robots.
- In our hypothetical, we assumed that the robots would need to build factories from scratch. In the real world, there will be a transition period where existing physical capital is reallocated to producing more robots.²⁸ This will help us to build robots faster – and by the time it's

27 In the actual world, this constraint will only bind once there are already more robots than can be produced by current stocks of physical capital. Before then, growth could be sustained by redirecting an increasing fraction of the world's existing physical capital towards robot construction.

28 See [here](#) for a sketch of how this reallocation might work.

necessary to build new factories, robot technology will probably already have improved a lot, which will lower the cost of building new factories.

- There is historical precedent for in-demand physical goods doubling every [~2 years](#), including solar panels, smartphones and electric cars.

Overall, we can tentatively say that initial robot doubling times are likely to be on the order of a few years, rather than months or decades.

How quick might robot doublings become by the time we reach the earth's carrying capacity?

This calculation has three steps:

1. How fast would robot doubling times be with current technology?
2. How many orders of magnitude will we scale up robot production before reaching the earth's carrying capacity?
3. How much will this scale-up reduce the doubling time, based on experience curves for cost vs production?

Step 1. [Above](#), we estimated that with current physical technology and abundant AI cognitive labour, robot doubling times might be about **one year**.

Step 2. Today fewer than 100,000 humanoid robots have been produced.²⁹ We expect that the earth's robot carrying capacity will be constrained by energy not by raw materials.³⁰ Solar energy hitting the earth is [2e17 W](#), whereas the human body uses [100W](#). If 5% of solar energy is used to run humanoid robots with efficiency matching humans, you could run $1e16/100 = 1e14$ humanoid robots.

That's a scale up of robot production of 9 orders of magnitude ($1e14/1e5 = 1e9$).

Step 3. [Above](#) we estimated that we might have to scale up robot production by 1-5 orders of magnitude to reduce the doubling time by one order of magnitude.

Conservative calculation: robot doubling times fall by $9 / 5 = -2$ orders of magnitude to **a few days**.

Median calculation: robot doubling times fall by $9 / 3 = 3$ orders of magnitude, to **a few hours**.

Aggressive calculation: robot doubling times fall by $9 / 1 = 9$ orders of magnitude to **less than a second**.

29 SoftBank's Pepper robot seems to be the largest-scale humanoid to date, with only [~27,000 units ever made](#). SoftBank's Nao robot follows at just [~19,000](#) deployed. The latest World Robotics – Service Robots survey (IFR, Oct 2023) counted [53,000 “social-interaction & companion” robots](#) shipped during 2022—this likely includes many of the robots already counted. Counts for recent humanoid robots like those from [Tesla](#) are much lower. Adding all of these worst-case non-overlapping figures ($27\,000 + 19\,000 + 100 + 53\,000$) still yields $< 100\,000$ humanoids produced in total.

30 We don't expect that raw materials will place a strong limit on growth, see footnote 15. By contrast, there seems to be a strong argument for energy limits. If earth produces significantly more energy than it currently receives from the sun, it will heat up significantly and become uninhabitable.

This suggests we could reach the doubling times of days or weeks suggested by the biological anchors.

Caveat: One big uncertainty in this calculation is that it does not consider the other types of physical capital (e.g. factories, machines, infrastructure). If some type of physical capital has a less favourable experience curve (and there's no alternative with a more favourable experience curve), then this could bottleneck growth and increase the doubling time.