

Comments from reviewer

Many thanks for sharing your thought provoking report! I've finally found the time to read over it and think about it. I wound up spending about 7 hours engaging with the material (reading the report and some background information, thinking about it, writing this email, etc.).

I found your analysis quite interesting and think your assessment based on the existing evidence and some speculation is quite reasonable. In sum, I think I agree with you that there is positive probability for AI fueled explosive growth, at least for some time. Whether and how long such growth could/would be sustained (even with AI) is more questionable.

Here are a few thoughts/comments I have:

--- To me, Jones style endogenous growth models with "small" scale effects (i.e. diminishing returns to "ideas production") are the most plausible explanation of the past growth experience (including the switch from Malthusian growth to exponential growth around the industrial revolution). It should also be straight forward to generate explosive growth via AI in these models.

For example, suppose that "ideas production" is the only possible source for explosive growth and is akin to "fishing a finite but large number of ideas out of a pond". This world would exhibit diminishing returns to ideas in ideas production and would be consistent with the current 2% frontier country growth. I find it quite plausible that AI could replace humans completely in ideas production, while humans would still be "necessary" in order to produce objects (including new "high-tech" gadgets like flying cars). If AI is sufficiently "intelligent" to keep catching fish at a constant and high rate, then this could surely generate growth above 2%. Will it be >30%? Who knows. But I'm with the ignorants on this one: why not?

--- If 30% growth is feasible, how long would it possibly be sustained? Let's for a moment become a bit more utopian and assume AI replaces humans in all production tasks. And let's further assume that part of the new "ideas" make explosive growth feasible. Let's also assume that humans are still "in charge" and own all the output and that, ultimately, the only incentive to further increase per capita GWP is human consumption. It seems like space and time are the only constraints in this world. Even with infinite amounts of output, humans only have finite lifetimes to consume finite amounts. Moreover, unless we move to Mars, there's only so much space on earth.

An interesting benchmark question is then: if Jeff Bezos and Bill Gates are currently defining the "maximum possible consumption", how many years of explosive growth would it take to bring all humans to this maximum consumption level? Then, how long would it take for all the empty space on earth to fill up with humans and bring all the new humans also up to that maximum

consumption level (given time and biological constraints)? Given that nobody has to break a sweat in this world, it is conceivable that nobody will complain about evenly sharing the pie, as long as per capita consumption is high enough.

Bill Gate's net worth increased by roughly \$4 billion from 2017-2018

(<https://www.thestreet.com/personal-finance/bill-gates-net-worth-14564023>). Let's assume that's a rough ballpark for an upper limit to annual consumption. According to the Penn World Table (PWT), average per capita GWP is \$23037.88 in 2017 USD. At worldwide 30% growth it would take approximately 46 years to bring average per capita GWP to \$4 billion. Estimates of the "carrying capacity" of the earth are currently 15 billion, so roughly double the current world population. So if population grew to carrying capacity, we would still reach \$4 billion per capita GWP in less than 50 years at 30% growth. Even if carrying capacity were 100 billion, it would still only take around 56 years.

That's surprisingly fast. These simple calculations would suggest that explosive growth could bring all humans on earth to an extremely high standard of living (as best as we can tell today) within roughly half a century. In contrast, at 2% growth it would take between 600-800 years to achieve this.

The reason why I bring this up: once we reach Phelps' "Golden Rule consumption" in this world, what is the point of further AI driven research? So I think this "capacity" argument is another reason why these very long run super-exponential reduced form projections don't seem super plausible to me. But explosive growth for a few decades seems conceivable.

--- There's another important question, which Anton has done some very interesting work on: What makes us believe that humans will stay "in charge". If it's AI that drives the explosive growth through ideas production, then when will AI have the idea to ignore what humans want? And at that point, how do we know whether AI would try to find ideas that increase growth, rather than other things that they "care" about?

--- In light of my back of the envelope calculations above, another interesting question is the incentive for humans/governments to "steer" innovation. For example, several people (myself and Anton included) have been thinking about the possibility of steering innovation in a way that focuses on complementing human tasks/skills rather than replacing them. But suppose we actually believed that explosive growth is feasible by replacing all humans in production with AI driven machines, and suppose we also believed my little back of the envelope calculation, that this could lead to worldwide "Golden Rule" consumption within a few decades. Then why would we even bother with redirecting AI innovation from labor-saving to labor-enhancing? These are of course some very big IFs. But I think this highlights why the question you're asking is a very important one from a policy perspective.

Reply from author

Thanks so much for these comments!

I really like the idea of using 'Bill Gates consumption' to think through how long we could in practice sustain 30% growth for a human population.

I had a couple of questions in response to your comments:

The reason why I bring this up: once we reach Phelps' "Golden Rule consumption" in this world, what is the point of further AI driven research?

I guess the reason would be to improve technology (and thus consumption) even further, at fixed resource consumption? It seems possible to me that significant technological progress might allow us to sustain 30% growth for longer than your calculation suggests.

For example, several people (myself and Anton included) have been thinking about the possibility of steering innovation in a way that focuses on complementing human tasks/skills rather than replacing them.

Interesting - is the motivation here to avoid disruptive effects on unemployment and inequality? Even if explosive growth is ultimately possible, I can imagine this would be a good idea on balance.

Reply from reviewer

Just a few quick responses:

*I guess the reason [behind further AI driven research] would be to improve technology (and thus consumption) even further, at fixed resource consumption?
It seems possible to me that significant technological progress might allow us to sustain 30% growth for longer than your calculation suggests.*

I agree. There are many reasons for why growth might continue. I just wanted to mention that (like in a Malthusian world) there may be some sort of "capacity" constraint. Here, my thought was "how much can humans possibly consume?". As long as more consumption is always better, then creating more output per person is a natural objective. But if the world is "satiated" at some point, then what is the new objective? Say the objective becomes "more variety" or something like that. Then, we would still want innovation but we might not see growth in GWP

per person. So “growth” in output might slow or stagnate, though innovation and growth in “happiness” might not.

As my little back of the envelope calculation suggests: with 2% growth, this is sort of an irrelevant question for the next 600-800 years. But with 30% growth, it may become relevant pretty soon. This is again something that, I think, highlights why thinking about your main question (can AI bring 30% growth?) is important.

For example, several people (myself and Anton included) have been thinking about the possibility of steering innovation in a way that focuses on complementing human tasks/skills rather than replacing them.

Interesting - is the motivation here to avoid disruptive effects on unemployment and inequality? Even if explosive growth is ultimately possible, I can imagine this would be a good idea on balance.

Yes. The idea is that, historically, the overwhelming private incentive for innovation appears to be cutting labor cost. But this may not be the social optimum because, for example, it may cause income and wealth inequality, which may have a variety of negative impacts on society (including the possibility of social unrest, coups, etc.). So from a cost-benefit perspective of a government, the question could be: should I invest in steering the technology in a way to avoid income/wealth inequality or should I battle the inequality ex-post via re-distribution. But this example takes as given that using humans in production is preferred to not using them. However, if the “no-humans” equilibrium could bring everybody to “maximum consumption” in a few decades, then it may not be worth the “effort” to figure out these other alternatives that need heavy “steering” but instead keep going with the re-distribution during the “transition”. But all of this depends on the confidence we have in making it to the “Golden Rule equilibrium” within a few decades. That’s what I meant when I said this highlights the importance of the question you are asking: how likely is it that AI could produce explosive growth soon?