General impressions

- What are your main thoughts about the report?

Missing considerations?

- Were the main considerations bearing on this question covered in the report? If not, what considerations were missing?

I would argue that you are missing two major considerations that bear on your conclusions. They are kind of related, but I think it makes sense to break them up.

The distinction between GDP/output and productivity.

The report is framed around the possibility of explosive growth in GWP, a measure of output. The issue is that much of the argument regarding the plausibility of explosive growth is made referring to models of R&D/idea production that related to *productivity* growth. Normally we might think that since productivity growth is so key to growth in GDP/GWP showing explosive growth in one amounts to showing explosive growth in the other.

Except that if we are talking about truly explosive productivity growth, we have to take into consideration the consequences for the provision of other inputs to production. Over history we've seen that our provision on inputs responds to the growth rate of output and to the level of output. The most obvious example here is the secular decline in working hours as countries become rich. We work far less per week, and far less per year, than our grandparents did. Many nations are shifting formally to shorter work weeks or shorter work years. The most plausible explanation for that, thanks to higher productivity, we can meet many of our material needs without putting in so much effort.

Extend that out to a situation where productivity growth is truly explosive. What's the response of workers when productivity jumps 30% per year? Even if only a third of that productivity growth translates to real wages, that's 10% growth per year in real wages, which is monumental. That's doubling your real living standards every 7 years. If you start working at age 21, by 35 your living standard would be 4 times higher, without even considering the additional earnings you might get from experience/promotion/etc.. To the extent that our young worker can invest some of their earnings in capital, that will be earning returns – given explosive growth – that are boosted by higher productivity growth. By 35 there is no reason this person cannot retire.

Alternatively, our 21 year old can decide that it would be boring to retire at 35, but that they don't want to waste *all* their time working during their life. So maybe they work 15 hours per week, or work one year and then go backpacking in Tibet the next. Because of the high productivity of the economy, this limited work effort is enough to afford all the basic material needs of life.

What consequence does this have on GWP? It might not go up. At all. If productivity goes up by 30%, but as a consequence the provision of other inputs like capital and labor *fall* by 30% for the reasons given above, then GWP will not change. Explosive growth in *productivity* does not imply explosive growth in *output*. We may well take advantage of productivity growth to slash our input provision, work the 15-hour weeks that Keynes speculated about, or do things like retire early or take extended time off. Which would be great! It doesn't mean explosive growth in productivity isn't worth it. It just means you have to be careful about framing this in terms of output.

Demand matters.

What happens to our selection of goods and services when productivity growth runs at 20-30% per year? History suggests that people tend to view many goods and services as complements. Yes, within specific sub-groups (e.g. shoes) different versions are close substitutes, but across those groups (e.g. shoes and live concerts) people treat them as complements and would like to consume some of both.

What does that do to the predictions of explosive growth? It suggests that it may "eat itself". AI or whatever will deliver productivity growth to some products faster than others, barring a truly unbelievable assumption that it can innovate *precisely* equally across every product in existence. When productivity grows more rapidly in product A than in product B (50% versus 10%, say), the relative price of product A falls relative to product B. Taking A and B as complements, what happens to the total expenditure on A (price times quantity)? It falls. We can get all the A we want for very cheap, and because we like both A and B, we have a limit on how much A we want. So total spending on A falls.

But growth in aggregate productivity (and in GWP, leaving aside my comments on inputs above) is a weighted average of productivity growth in all products. The weights are the expenditure shares. So in the A/B example, as A gets more and more productive relative to B, the productivity growth rate *falls* towards the 10% of product B. In general, the growth rate of productivity is going to get driven towards the *lowest* productivity growth rate across the range of products we consume.

And the faster that productivity grows in product A, the sooner the aggregate growth rate will fall to the productivity growth rate of B. So a massive question for this report is how widespread explosive growth is expected to be. Productivity growth in *all* products of 10% forever would deliver 10% growth in productivity forever (and perhaps in GWP). Great. But productivity growth of 100% in A and 0% in B will devolve into productivity growth of 0% over time.
This has nothing to do with the nature of R&D or the knife-edge conditions on growth models. This is simply about the nature of demand for products.

Why might this be wrong? Well, it may be that AI can deliver productivity growth at high rates for every single product. If the minimum growth rate in productivity for a product is like 20%, then the economy will tend towards that minimum.

The other issue would be that perhaps products are not complements the way I described. But as I said, all the evidence we have over time suggests that they are. You can see this in the decline of manufacturing as a share of output, even though productivity growth in manufacturing is way higher than in services. We have gotten very productive at making refrigerators over time, and most people take advantage of that to … take a vacation or get surgery or take a class. People don't accumulate 10-12 refrigerators.

The last "out" here would be new products. To the extent that AI can deliver a brand new category or type of product, that would prevent some of this from happening. If AI invents good C, then a bunch of expenditure will flow into spending on C, and at least for a while that could keep growth in output high.

Reasons to expect explosive growth

- The report identifies the key potential driver of explosive growth as "increasing returns to output-bottlenecked inputs". Do you agree?

The report identifies three main reasons you might expect explosive growth:

1. **Explosive growth models.** These seem to be a good fit to long-run data and their predictions might be trustworthy if AI substitutes effectively for human labour.
2. **AI-specific models.** Simple extensions of standard economics models suggest explosive growth could happen if AI allows capital to substitute more effectively with labour
3. **An ignorance perspective.** Growth has increased in the past; we don't understand why; maybe growth will increase again.

- Do you think these are good reasons to expect explosive growth? Are there others we've missed? Would you carve up the space of arguments differently?

  I'm on board with your argument that we cannot dismiss explosive growth simply because we didn't observe it in the past. As you very rightly say, asking someone in 1750 what the possibilities for growth were in the future, they'd say 0% per year, given all their available evidence. So just as I argued above that starting with extrapolation may not be the best framing, that doesn't entitle me to just extrapolate growth rates of 1-4% without some justification. But that doesn't really constitute an argument *for* explosive growth.

  The single best argument for explosive growth that I think you have is the substitution of AI/automation into the R&D process. As you note, all of our standard models of R&D and productivity growth depend on this feedback loop where more people mean more ideas. We've disciplined those models with observed data from the 20th century to add in this dampening effect that means higher productivity slows down productivity growth. We essentially say that more ideas make it harder to find the *next* idea. And you might look

at some work by Ben Jones at Northwestern on this, as he has some interesting papers on how innovation and invention have become more difficult over time as researchers have to get up to speed on a larger body of existing knowledge.

Anyway, how could AI break us out of that? Basically by removing the limit placed by the growth in the number of people. We assume that the number of ideas is proportional to the number of people, and thus the growth rate of people puts a cap on the growth of ideas. But if AI is generating the ideas, and moreover the *capability* of AI is growing over time, then the cap placed by people doesn't apply any more. If the "stock" of AI grows at 20% per year, but the stock of R&D workers can only grow at 2% per year, then the growth rate of productivity should go up by a factor of about 10 if we plug AI into the R&D production system.

Note that this doesn't require us to dismiss the idea that higher productivity makes R&D harder. Even if R&D gets harder over time because we learn more and more, that doesn't stop rapid growth in AI from allowing for rapid growth in productivity. The key element here is the speed of the input to R&D. Since AI can grow, conceptually, much faster than the number of people working R&D, that can stimulate much higher productivity growth.

Of course, my caveats about translating productivity growth into output growth, and demand effects, still apply. But if there is a possibility of explosive growth, to me it comes from the possibility of massive growth in R&D effort.

Explosive growth, conditional upon AGI

- Would you expect explosive growth to occur, conditional on us developing AI systems that can replace humans in all tasks?
    - What would be your (unavoidably speculative) probability on explosive growth, conditional on developing AI systems that can replace humans in all tasks?

All that said, I think the probability of explosive growth in GWP is very low. Like 0% low. I think those issues I raised above regarding output and demand will bind and bite very hard if productivity grows that fast. I will steal this line from Alex Tabarrok. "The growth rate in heaven is zero". People will respond to massive changes in productivity, and those changes are very likely to offset a lot of the growth in GWP that you are banking on.

That's different than the probability of explosive growth in productivity. That's probably like 20%? I still have a somewhat low expectation on that. In part that is because I think we over-estimate how many things AI could potentially innovate on. It's great for AI to innovate a new lighter material that makes airplanes less expensive, for example. But how does AI innovate on a massage parlor, a restaurant meal, or an exercise class? More to the point, I think there will come a point where people don't *want* an

innovation. I personally don't want my meal at a nice restaurant with my wife to be more efficient or faster.

My own pet prediction is that if/when AI provides more rapid productivity growth in some areas, you will see a very deliberate shift of people to providing goods and services in very "traditional" or "by-hand" methods. They'll be sold specifically on the basis of their inefficiency. That's pure speculation, but I think you can see hints of it already.

Knife-edge conditions

The report claims that it is difficult to find endogenous growth models that produce exponential growth without knife-edge conditions, given that the population is expected to grow sub-exponentially. It takes this to be a theoretical reason to doubt that growth across the rest of this century will be exponential.

- Do you agree? Do you find the reasoning of the report convincing on this topic?

I guess I'd question this characterization of endogenous growth models. There *are* growth models with very knife-edge conditions. Paul Romer's original model was one, and it had counter-factual predictions as it turned out. And yes a bunch of the early wave of endogenous growth models had this same issue. The earlier models (AK, etc..) also had very particular assumptions necessary that didn't make sense. Yes, but I don't think any of these models are really used as the basis for study any more?

The "semi-endogenous" growth models that Chad Jones is kind of the champion of don't depend on a specific parameter value to generate growth. They rule out the knife-edge conditions of the earlier generation because, at least so far, they don't match the data.

I think you bring up that the semi-endogenous models depend on the "knife-edge" condition of exponential growth in population. I guess I don't see that as some kind of knife-edge, the way that earlier models need an (unobservable) parameter to be "just so" to make things work. The semi-endogenous models imply that growth in productivity is proportional to the number of people (really, the number of people doing research). Whether that number grows exponentially or not is immaterial to that finding.

What that implies is that yes, if the population growth rate falls (as we expect or observe) then this should pull down the growth rate of productivity (and probably GWP). *So far*, that would be consistent with how we observe growth working. To your point, changing the relationship of productivity growth and population growth by adding AI might alter that prediction.

Other

- Any other comments on the report, or on the process as a whole?

The one other area I might suggest you take a look at is some old stuff by Solow and Olivier de La Grandville. In particular, they looked at how the elasticity of substitution between labor and

capital could generate endogenous growth, if that elasticity was high enough. That is, if capital could keep acting in place of labor, then you could get output per person to grow just by acquiring more of that capital. They laid out some conditions under which the elasticity was big enough to do that.

What I'm thinking here is that AI might in some sense act as a way of increasing that elasticity, allowing us to jump to this endogenous growth situation? Say that today it is hard to substitute capital for some labor. But AI, perhaps in the form of a robot, makes I plausible to substitute capital for labor.

It isn't productivity growth, per se, but achieves possibly the same result.

Framing of the question

- The report is focused on assessing the plausibility of 'explosive growth', defined as >=30% growth, occurring by 2100. Does this framing make sense to you? Do you think we should have framed it in some other way?

  I do wonder about the framing in the sense of the "null". From the outset the "null" is that explosive growth is expected based on a simple extrapolation from the past (e.g. growth rates have generally increased century by century). So the report takes that as the baseline, and is in essence examining the evidence/arguments that might cast doubt on that extrapolation.

  A very different "null" is that growth rates of GWP will continue to plod along in the 1-4% range we've seen since, say, 1850. Then the onus is on the report to show that explosive growth is plausible and cast doubts on the idea that we will plod along. I think this framing is more sensible. Why? Because for all the knife-edge issues that might be true about our theories of growth, they are built on plausible mechanisms, and it is *very* hard to get them to deliver explosive growth. The extrapolation is theory-less, and so I don't see any reason to take it more seriously.

  Here's my analogy to explain what I mean. When my daughter was born, the growth rate of her weight in the first year was about 200% - she tripled in weight. Were I to extrapolate that growth rate, I would predict she would have been about 2.7 billion pounds by the time she was 18. She is not 2.7 billion pounds, as it turns out. She's barely over 100.

  There is plenty of decent theory about why her growth rate slowed down. And if you kind of muck around with the biology of it, it's also pretty knife-edge. There are a lot of thresholds of hormone levels and such that work "just so" to ensure our bodies level off at the levels they do. Despite that, the theory of body development is a way better guide to future weight growth than extrapolation.

So I think the report might benefit from a “Here’s the conditions under which explosive growth *could* occur” as opposed to “Here’s why explosive growth *will* occur”.

Permissions

- Would you be willing for us to publish your answers to the above questions alongside the report?

  Sure!

- Would it be OK for us to publish your name alongside your comments?

  Sure!

Email exchange

I’ve grouped the discussion by topic to make it easier to follow.

Framing of the question

Author
Re framing, I had actually intended the report to have something like the framing you recommended. I.e. I didn't mean to set things up such that explosive growth was a 'default', but rather to contrast two possible empirical 'defaults': the steady growth over the last 150 years vs the increasing growth over the past 10,000 years). If you happened to have any quick thoughts on how I could make my intended framing clearer I'd be keen to make those changes. (I've just edited the abstract which I think might have been to blame!)

Reviewer
Tom, on the framing it’s not over the top, but as I read I kind of got the impression of working with the “null”, as I said, that growth would be that high.

I think I would set up the growth econ “norm” of modern sustained growth of around 2% per year as more of a straw man to open it? Something like “For over a century the growth rate of GDP per capita has averaged around 2% in developed parts of the world, and this fact has informed most research on economic growth in that era. Looking forward, most models of R&D and innovation would predict continued growth at those rates. But is this predication warranted? Rates of growth have accelerated over the course of history, and the opportunities presented by AI and rapid automation may eliminate the constraints on growth that those models take for granted.”

I think something like that sets it up that we have a baseline understanding of modern growth, but you (rightly) can ask whether it is the appropriate model for the future.

Demand matters

Author
I've been thinking a bit about the argument that demand considerations like the ones you raise might prevent explosive growth from occurring. I buy that it would push in the direction of making 30% GWP growth less likely, but currently don't think it rules it out.

Here are some of my thoughts:

1. Average global average income is currently ~1/6 of US average income. If AI allowed the world to rapidly catch up to the global frontier, that would involve a ~6X GWP increase and potentially very rapid growth. It doesn't seem like demand considerations would rule this out.
2. Further, there are people alive today who enjoy incomes 10X - 1000X higher than the US average. If AI allowed the world to catch up to *these* incomes, that would involve a 60 - 6000X GWP increase and again potentially very rapid growth. Again, it doesn't seem like demand considerations would prevent this from happening?
3. In competitive scenarios, the demand for a particular product might not decrease even when we get much better at producing it. E.g. if AI allowed military power to grow extremely rapidly, the 'products are complementary' model predicts we'd just spend less on military and focus on other harder-to-improve things. But in reality countries trying to increase their relative influence might spend *more* on military in this situation.

I know that's a mouthful, but would be interested in any thoughts you have!

Reviewer
In general I think that yes, the idea of income distribution (especially world-wide) is a good counter-argument to my demand-side argument. That is, even if I'm right about how demand works, there may be so many people out there at low levels of living standards that the effects of demand really will be small in quantitative terms as billions of people build up the durable goods and living standards of advanced economies. On the other hand, if GWP really were to grow at 30%, it wouldn't take long before everyone was rich enough for those demand effects to kick in. But to your point, maybe that still is 100 years away.

I don't think using comparisons to rich people today helps you in this. Yes, Jeff Bezos has roughly 60,000 times my wealth, and I'm sure his yearly income is a similar order of magnitude higher than mine. But I guarantee you that Bezos does not have 60,000 times as many refrigerators as I do. Or cars. Like, maybe he has houses, in total, worth 60,000 times as much as mine.

I think Bezos makes my point, in many ways. The vast majority of Bezos' income does not get spent on any good or service, it gets plowed back into investments (not part of GDP) that grow his wealth even further. Beyond that, what he *does* spend on is almost certainly a very low productivity basket of goods and services. What I mean by that is that Bezos is probably consuming services or goods that are, by design, low productivity. A bespoke suit made from yarn hand-spun by Peruvian villagers is probably *amazing* in terms of comfort and fit. It is also a disaster in terms of labor productivity.

On your last point (defense), you are right. To the extent that goods are seen as substitutes, then demand works in your favor. If you are willing to buy *more* of the cheaper goods, then eventually your economy is dominated by the high-productivity product and productivity growth approaches the maximum.

Author
I agree with your point that Bezos spends his money on ultra-low productivity goods. To check my understanding, are you saying that if everyone consumed the same basket of goods as Bezos currently consumes, average consumption wouldn't really have risen by 10,000X? Even though Bezos currently pays 10,000X for his consumption compared to the average US person (let's assume), if everyone consumed that same basket of goods then its price would fall and people's real consumption would rise by less than 10,000X. Am I understanding that correctly?

I had one last question relating to the demand point. The question relates to a scenario where AI and robotics advance to a stage where we can set up self-replicating robots: robots that can autonomously mine the raw materials and build the factories needed to make more robots. Without human involvement, the number of robots and AIs can grow exponentially (until resources run out). Let's also assume (to your point about demand) that there are some products (e.g. haircuts, massages, priests) where these robots can't directly provide value, and others (e.g. butlers) where additional robots just get in the way. In this scenario, one prediction would be that we would create enough AIs and robots to serve our immediate needs, and then produce no more. But another prediction would be that we'd make a huge number of robots and AIs because we could then use them to do R&D into improving human health (for which there is a very large willingness to spend) and other areas. In other words, we'd make loads of AIs and robots not for their direct effect on the economy, but because of the R&D they could do. (Another reason to create many robots in this scenario would be for military advantage, as we've already discussed.) In this scenario, do you think that demand considerations would prevent humans creating huge numbers of robots and AIs? Of course, this is inevitably a speculative question! I'm just interested in how powerful the demand effect might be in preventing growth from accelerating if the technology was sufficiently advanced.

Reviewer
On the Bezos point, I think it helps to not think in monetary terms (10,000X) because that isn't really what we're after. For productivity (whether GDP/worker or TFP) it's really about real goods per unit of input. Bezos gets relatively few goods per unit of input (the hand-spun bespoke suit I

mentioned). If we get to the point where everyone can get one of these suits then there are two things going on:

1. It must be that our productivity is really high, because the only reason you buy hand-spun bespoke suits is because your food, housing, shelter, etc.. are easily taken care of.

2. People are taking advantage of that high productivity to buy very low productivity things like the bespoke suits. Meaning that productivity isn't going to grow very fast.

So I think my answer is that no, overall productivity wouldn't quite scale up by as much as it seems if we all were Bezos. As productivity rises, people use that switch consumption to low productivity goods, slowing down productivity growth. So maybe the way to say this is that yes, for sure productivity can become much higher, but higher productivity will tend to put a drag on further productivity growth. And the open question is how strong that drag is. If productivity doubles does that put enough drag on it that we "max out" at productivity 2X todays level? Or will we max out at 10X, or 100X, or 1000X? I'm sure it will max out, but I'm open to the possibility that it's only at 100 or 1000X, and we've got decades or centuries before it does.

On the robots, I think you are getting at the right consideration about demand. What makes cancer treatments or other health innovations potentially different is that you can imagine that they do not have this satiation point (unlike refrigerators or cars or other material things). People will probably always want to improve their health even by miniscule amounts. You might even argue that health innovations have this accelerating demand where if the material conditions get better your demand for health goes up even *more* because now it really pays to be healthy and alive for as long as possible.

So in your robot/AI scenario, if we were producing them to generate new ways to make ourselves healthy and live longer, yes, I could see that we'd continue to produce them (or let them produce themselves ad infinitum).

And that kind of gets at why demand matters. The difference between health improvements and butlers is that I always want more health, but I won't always want more butlers.

Author
On the Bezos point, I still feel a little confused. On way of defining someone's consumption is (number of goods they consume) * (average price per good). On this definition, Bezos' consumption is currently 10,000X the global average (so I've assumed). If everyone consumed the same basket of goods as Bezos currently consumes, and prices stay constant, then it seems like the global average consumption will now be the same as Bezos *current* consumption. I.e. global average consumption will have increased 10,000X. And I'm not sure exactly where this argument goes wrong?

Is your point that we shouldn't care about consumption as measured by monetary value, but only about consumption as measured by the number of real goods consumed? You might grant that global average consumption could rise 10,000X, measuring consumption by monetary value, but deny that it could rise 10,000X if we instead measure consumption by the number of real goods consumed. As you say, Bezos purchases low productivity goods and so he doesn't actually consume 10,000X as many real goods as the global average today.

Reviewer
I think we have to distinguish between productivity and the just raw consumption. Let's say Bezos consumes 10000X more real goods than the average person. So if everyone consumed his basket of goods (somehow) real consumption would be 10000X higher on average. But that doesn't mean productivity would be 10000X higher. It might be lower (probably not, but hypothetically yes).

Right now, let's say there are 99 people and 1 Bezos. The 99 of us each consume 1 real unit of goods, and that real unit takes 1 unit of inputs to make. Bezos consumes 10000 real units of goods, and each of those takes 2 inputs to make. We produce 99x1 + 1x10000 = 10099 units of goods, using 99x1 + 2x10000 = 20099 inputs. Productivity is 10099/20099 which is bigger than 1/2.

If we all consume the same amount as Bezos, then we produce 100x10000 = 1,000,000 (definitely more than before). We use 2x1000000 = 2,000,000 inputs. Productivity is exactly 1/2. Productivity went *down* because now we all consume low productivity items (but a lot of them). You can make the difference in numbers more dramatic by having more of "us" relative to Bezos.

So in some sense the question is what are we trying to measure, or what counts? If all we care about is raw consumption, then it's better that we all consume what Bezos consumes. And to some extent that makes sense. I think we'd all agree the world would be better if most people (who are really poor) could consume more things like food, clothing, and shelter.

But at some point do we care about productivity, especially in terms of how many inputs/resources we need to produce those consumption goods? Do we want to use up 2,000,000 units of inputs to ensure everyone can get their bespoke llama-hair suit? Probably not.

Now, if productivity increases in producing both kinds of goods, we could really get somewhere. If the "Bezos bundle" dropped to 1 unit of input per unit of output, that would be great. We could consume more and have higher productivity (but might still worry about total input use). My worry is that it might not be possible to increase the productivity of the "Bezos bundle" because in large part the whole *point* of the Bezos bundle is that it is unproductive. You can't have a bespoke suit that is mass-produced.

The distinction between GDP and productivity

Author
One thought I had in regards to productivity vs GWP: if AI ultimately automates (much of) goods production (as well as R&D) then humans could work fewer hours without inputs to goods production declining. In which case the gap between productivity growth and GWP growth would shrink. Does that sound right to you?

Reviewer
In terms of your question on productivity and GWP, I see what you are getting at. If you shed inputs, then at some point Y = A, and productivity is all that matters, meaning their growth rates are identical. In principle I can see that. In practice I don't think that would be a plausible outcome in the next 50-100 years. The reason is that while I think it is plausible for us to shed a lot of labor out of production, I don't think we will shed capital nearly as fast. But that is just my own speculation on the pace of these things.

Author
Re productivity, I think that what I had in mind is more like replacing human workers with AI controlled robots, so that at some point Y = AK. To have this we would need to shed human labour, but wouldn't need to shed capital. Does that make sense?

Reviewer
Got it. I think that Y = AK is probably eventually right - and I think you might have even mentioned the Seater and Peretto papers that think about this? But even in that case, at some point if productivity goes up we might choose to lower the amount of capital we use rather than keep accumulating it. How many robot servants can one family use?

Author
That makes sense - thanks! So this again comes down to the question of demand? One interpretation of what you're saying is that even if we could have 30% growth of output without working, we might prefer to have less capital and less output growth.

Reviewer
Yeah, I think that's right, demand matters at some point. (By the way, not "Aggregate demand" in the business cycle sense, just plain old "Preferences for things"). In the same way that at some point you run out of reasons to buy another refrigerator, you might run out of reasons to buy another personal servant robot.