

Answering the Techno-Pessimists (complete)

 noahpinion.substack.com/p/answering-the-techno-pessimists-complete



Noah Smith, May 30, 2021

When I started this Substack six months ago, I made it explicitly a techno-optimist blog. A number of my earliest posts were gushing with optimism over the magical new technologies of cheap solar, cheap batteries, mRNA vaccines, and so on. But a blogger at a blog called Applied Divinity Studies wrote a post demanding more rigor to accompany my rosy projections, and putting forth a number of arguments in favor of continued stagnation. Heavily paraphrased, these were:

1. We've picked the low-hanging fruit of science
2. Productivity has been slowing down, why should it accelerate now?
3. Solar, batteries, and other green energy tech isn't for real
4. Life expectancy is stagnating

So I decided to write a series of posts addressing all of these arguments. Here's the whole series in one post.


[Subscribe now](#)

[Share](#)

Part 1: Life Expectancy

Is life expectancy stagnating?

The blogger at [Applied Divinity Studies](#) posted the following graph:



They write:


So sure, there is a slight uptick, but basically it is still at a plateau with growth far below historical levels.

If you believed in stagnation when the paper first came out, you had better continue believing it in now.

Tyler Cowen and Ben Southwood are not techno-pessimists, but they also made the argument that plateauing U.S. life expectancy is a sign of recent technological stagnation.

But there are two big problems with this argument. First of all, life expectancy hasn't actually stagnated. Here's a graph with some other developed nations included:

Life expectancy, 1950 to 2019



Source: Riley (2005), Clio Infra (2015), and UN Population Division (2019)

OurWorldInData.org/life-expectancy • CC BY

Note: Shown is period life expectancy at birth, the average number of years a newborn would live if the pattern of mortality in the given year were to stay the same throughout its life.

You can see that the stagnation is almost entirely a U.S.-specific phenomenon. The U.S. falls behind other advanced countries in life expectancy growth around 1980, and absolute life expectancy in the mid-80s, and it never catches back up. It's also possible that the UK has experienced a stagnation in life expectancy just in the past few years, but too early to tell. The U.S. is really the only one who has clearly plateaued.

Presumably, the U.S. and other developed countries have access to similar life-extending technologies. Thus, the one-country life expectancy stagnation in the U.S. is almost certainly due to non-technological factors — lifestyle differences, differences in the health care system, etc.


Life expectancy isn't a great measure of technology

First of all, here's a graph of GDP per capita in England, where they've been keeping fairly consistent records for a very long time:

GDP per capita in England

Adjusted for inflation and measured in British Pounds in 2013 prices

Our World
in Data



Source: Broadberry, Campbell, Klein, Overton, and van Leeuwen (2015) via Bank of England (2020)


Note: Data refers to England until 1700 and the UK from then onwards.

OurWorldInData.org/economic-growth • CC BY

You can see living standards have gone up by a factor of about 30 since the Middle Ages. That reflects massive technological progress. But during that time, life expectancy has increased only by a factor of two:

Life expectancy, 1543 to 2015

Our World
in Data



Source: Riley (2005), Clio Infra (2015), and UN Population Division (2019)

Note: Shown is period life expectancy at birth, the average number of years a newborn would live if the pattern of mortality in the given year were to stay the same throughout its life.

OurWorldInData.org/life-expectancy • CC BY

The reason is that much of the big increase in life expectancy was actually a drop in infant mortality (and to a lesser extent, child mortality). We all hear about how people used to only live to age 35 back in Ancient Rome, or whatever. But that's not right. Basically, if you made it past the danger zone, you could expect to live to age 60 or so.

Mortality rates of children over the last two millennia


Our World
in Data

What about the rise in adult life expectancy? Adults used to live to be 60, now they live to be 80. That's a big improvement!

5/25

Life expectancy at age 45, United States, 1900 to 2014

Our World
in Data

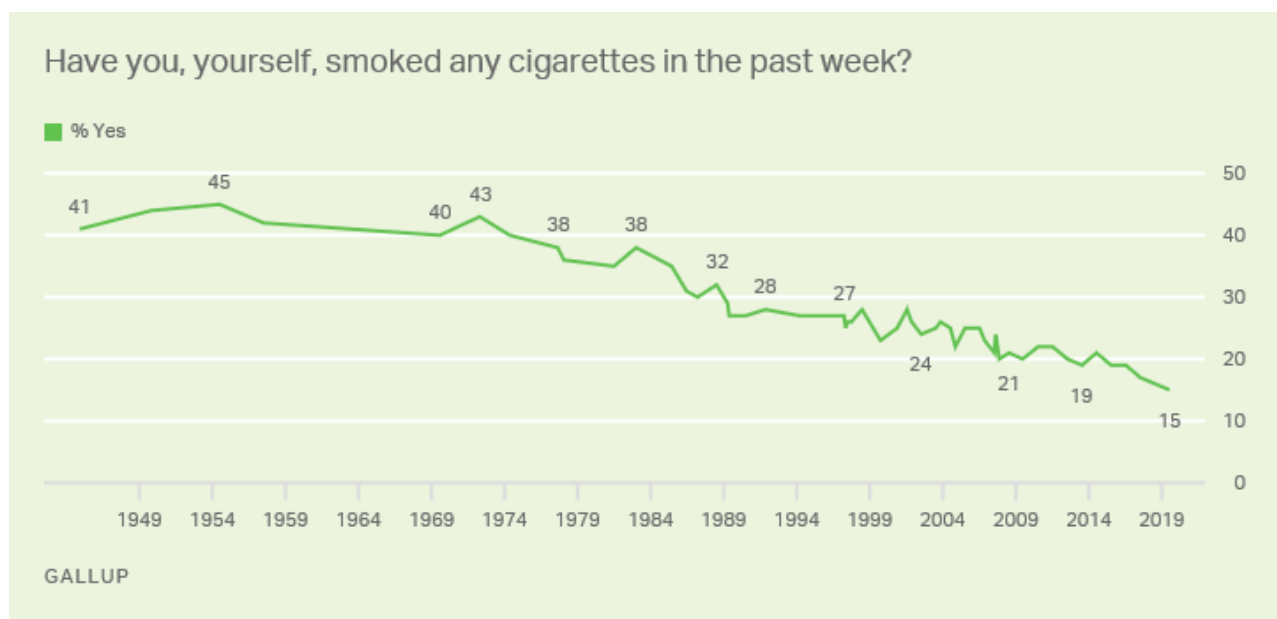


Source: OWID based on Hacker (2010) and the US Social Security Administration (2017)

CC BY

For American women, there was a big acceleration between 1930 and 1950 (probably due largely to a big one-time drop in maternal mortality), but otherwise a fairly steady increase that doesn't really match up with productivity or with our general sense of when the big, important inventions happened.

For men, the big acceleration in adult life expectancy starts around 1970 — basically the exact same point that productivity slowed down! A big driver is probably lifestyle changes, e.g. the drop in smoking:



Anyway, it's clear that although technology to reduce infant and maternal mortality was very important in giving life expectancy a one-time boost, in general life expectancy is not a good proxy for other broad measures of technological progress.

In conclusion: We might get life extension technology in the future, but we shouldn't expect a continued steady upward climb of life expectancy as overall technology improves.

Part 2: Solar and Batteries


Slowly, then all at once

The Applied Divinity Studies blogger posts a chart showing the amazing decline in solar power costs during the 2010s:

Solar PV module prices

Global average price of solar photovoltaic (PV) modules, measured in 2019 US\$ per Watt.

Our World
in Data



Source: LaFond et al. (2017) & IRENA Database

OurWorldInData.org/energy • CC BY

Observing this chart, the blogger asks: If solar cost declines are going to increase productivity growth in the 2020s, why didn't they do so in the 2010s?

Can you show me where The Great Stagnation was? Maybe there was one from 1990 to 2006, but you can't reasonably look at this graph and infer that 2006-2017 was a bad decade.

And yet somehow, the Stagnation Hypothesis was convincing in 2012 as the price dropped from \$1.68 \$/W to 0.88 \$/W. And it remained convincing in 2017 as they dropped further from 0.55 \$/W to 0.46 \$/W.

What exactly about the last year feels compelling in a way that previous progress was not?...How is this "Optimism for the 2020s" as Noah writes or a "Crack in the Great Stagnation" as Watney suggests if it's been going on for the last 10 years?

It's a good question. In fact, the decades since 1976 have ALL been amazing years for progress in solar technology (though the 90s and 2000s less so). But before we can start seeing that progress in the aggregate statistics, we need two things:

1. We need solar to be **installed at large scale**, and
2. We need solar to be not just cheaper than fossil fuels, but **substantially cheaper**.

Progress in solar technology has gone through several phases. [A 2018 paper](#) by Goksin Kavlaka, James McNerneya, and Jessika E. Trancik chronicles the first two of these. In the first phase — roughly, the mid-1970s through 2000 — solar power was mostly confined to the laboratory, since it was too expensive to use as a power source for industrial or commercial applications. During this time period, solar cost drops were driven by government research efforts.


Around 2001, things shifted, when solar got cheap enough that companies started to deploy it at a big enough scale for learning curves to start having an effect. As more solar was installed, companies learned how to do it more cheaply, and took advantage of more economies of scale. That kept the cost declines going. And as cost declined further, more solar was installed. This process is still going to this day.

As solar got cheaper, the amount of solar power installed in the world increased more-or-less exponentially (remember that an exponential is a straight line on a log scale):

Solar power generation

Electricity generation from solar, measured in terawatt-hours (TWh) per year.

Our World
in Data




Source: Our World in Data based on BP Statistical Review of World Energy & Ember

OurWorldInData.org/renewable-energy • CC BY

The third phase of solar was when it started to become an appreciable fraction of the world's electric power generation. This was sometime in the 2010s, depending on how large you think "appreciable" is:

Figure 2: Global power generation mix



Source: BloombergNEF

But notice that even as of 2019, solar was *still just a small fraction of total power generation* — about 3% of the global total. You can't make an impact in the aggregate productivity statistics by changing 3% of global electricity consumption to a different source. That is why the amazing technological progress in solar hasn't yet translated to faster *economic* progress. **Technology is embodied.** It doesn't automatically change economics just because you *invent* it; you also have to *build* it.

But now it's getting built. As many COVID denialists learned to their great dismay in 2020, exponential curves move fast. Bloomberg New Energy Finance's graph above is probably way too conservative in its prediction of how much of our electricity we'll get from solar, as Ramez Naam explains. The real growth rate of solar has tended to be more than 3 times as fast as BNEF forecast in the past.

An energy source that is 3% of electricity generation will not show up in the productivity statistics, but one that is 30% of electricity generation definitely will. And THAT is why solar cost drops didn't show up in aggregate productivity statistics in the 2010s, but will definitely show up in the future. (Note: For those who are concerned about intermittency, realize that we'll just overbuild for winter/clouds/morning/evening, and use batteries at night.)

Everything I've just said goes for batteries as well. Electric cars are just now hitting the point where they can compete with internal combustion cars in terms of usability and affordability, and the market for electric cars is just starting to take off. Notably, some

carmakers are starting to cancel internal combustion models in anticipation of the switch.

Like I said: Slowly, then all at once.

But in order for solar and batteries (and other green energy technologies like wind and hydrogen) to really make a big impact in the aggregate productivity statistics, they can't just replace fossil fuels — they need to be substantially better than fossil fuels ever were.


That looks likely to happen soon.

Cheap energy is what boosts productivity

Clean energy isn't just clean; it also has the potential to be very cheap. We've had hundreds of years to improve and refine fossil fuel technologies; they're not getting much cheaper. Fracking was an exception, but it's a pretty modest improvement compared to what's happening in renewables:

The price of electricity from new power plants

Electricity prices are expressed in 'levelized costs of energy' (LCOE). LCOE captures the cost of building the power plant itself as well as the ongoing costs for fuel and operating the power plant over its lifetime.



Data: Lazard Levelized Cost of Energy Analysis, Version 13.0

OurWorldinData.org – Research and data to make progress against the world's largest problems.

Licensed under CC-BY by the author Max Roser.

There's really no end in sight yet for improvements in solar and batteries. Cost drops are continuing simply from scaling up, and new materials and technologies are on the horizon that could generate continued price declines per unit of energy.

In other words, solar and batteries have a long way to go.

And that's important, because simply replacing fossil fuels with solar doesn't automatically improve productivity by a noticeable amount. If you replace a coal plant with a solar plant that generates energy for 97% of the coal plant's price, that's good for the environment, and it creates some jobs, but it doesn't really boost productivity much because the energy generated is only a tiny amount cheaper than before.

The real productivity gains happen when new energy technologies get *much* cheaper than old ones. If solar plants can produce electricity at half or a third of the (levelized) cost that we got with coal or gas plants, that's a big deal economically. The productivity slowdown of the 1970s happened in large part because oil stopped being cheap; likewise, much cheaper energy would enable a productivity boom.

In fact, there's a hypothesis that physical technologies are more effective than digital or other technologies at goosing the productivity numbers. I laid that hypothesis out in a previous post, which I will now quote from:

Another possibility is that “atoms” innovation, unlike “bits” innovation, unlocks the potential for more extensive growth; instead of simply making us more efficient or more creative at using resources, physical innovation may allow us to gobble up more of the world around us. (As Tyler Cowen and Ben Southwood put it, “many scientific advances work through enabling a greater supply of labor, capital, and land, and those advances will be undervalued by a TFP metric”). A 2013 retrospective on productivity growth by Robert Shackleton of the Congressional Budget Office suggested:

Finally, the sweep of the 20th century underlines the extent to which long-term TFP growth and economic growth in general have been influenced by the development of energy and transportation infrastructure suited to the expansion of suburbs.

Of course, this is a hypothesis; I can't prove things will work this way. But even if this hypothesis is wrong, energy is a very important input into pretty much all production processes, and so much cheaper energy will improve productivity growth — as soon as solar and batteries get much cheaper than fossil fuels and are installed at sufficient scale!

Sustainability is also real progress

So the two reasons solar and batteries haven't yet produced higher aggregate productivity growth are that they haven't been installed in large volumes yet, and they haven't gotten so cheap that they dramatically lower the price of energy yet. But current trends give us every reason to believe that both of these things will happen.

But in addition to measured productivity growth, green energy has another benefit: Sustainability. They help prevent catastrophic climate change, for one thing. And also, because the supply of sunlight is not limited, solar frees us from the need to constantly invent new and better extraction technologies just to hold energy costs constant.

Neither of these things adds to current productivity, but they both add to expected future productivity — because future productivity won't be nearly as high if we experience climate catastrophe and fuel keeps getting harder to extract. As Tyler Cowen writes in his book *Stubborn Attachments*, the future is very long, and if we value future humans, sustaining productivity over the long haul is more important than bumping it up for a moment. So sustainability really can be viewed as a form of progress, even if it doesn't juice the statistics in the present.

Part 3: Productivity Trends

Technological progress and TFP are not the same thing


Generally, the main evidence that technological progress has slowed down comes from something called total factor productivity (TFP). Many of the Applied Divinity Studies blogger's arguments rely on TFP as the fundamental measure of technology. So let me explain what that is.

Economists generally think that economic output (GDP) depends on two things: A) the amount of inputs like capital and labor, and B) technology, which turns those inputs into productive output. To measure TFP, you have to make an assumption about the functional form of the production function that transforms inputs into output. Usually we just assume something that looks like this:

$$Y = A \times K^{\alpha} \times L^{\beta}$$


Y is GDP, A is TFP, K is the economy's total amount of capital, L is the economy's total amount of labor, and α and β are just some numbers.

Using this method, we get something that looks like this:



This looks like mostly steady linear growth, punctuated by a couple of stagnations.

But we don't want linear growth! We want exponential growth! A constant linear increase means that the percentage growth rate is going down. To look at the exponential growth rate, we should look at the logarithm of productivity:



As you can see, each period of productivity growth now looks a bit slower than the last, meaning that the exponential growth rate is slowing down.

In fact, if you adjust for the changing utilization of capital and labor, as my PhD advisor Miles Kimball did in a paper with Susanto Basu and John Fernald, you find that the growth rate of “true” TFP has stagnated even more than the above graph would suggest. In other words, we’ve been using our machines and buildings and stuff more intensively, disguising some of the true TFP slowdown.

OK, so that’s the TFP slowdown. But here’s the thing: **TFP is not the same thing as technology**. The word “technology”, as we commonly understand it, includes stuff like computer chips, car engines, and procedures for making cement. Economists would broaden that definition to include things like business management techniques. But even with that broad definition, there’s plenty of stuff that can affect TFP that most of us would agree does not represent actual technology. For example:

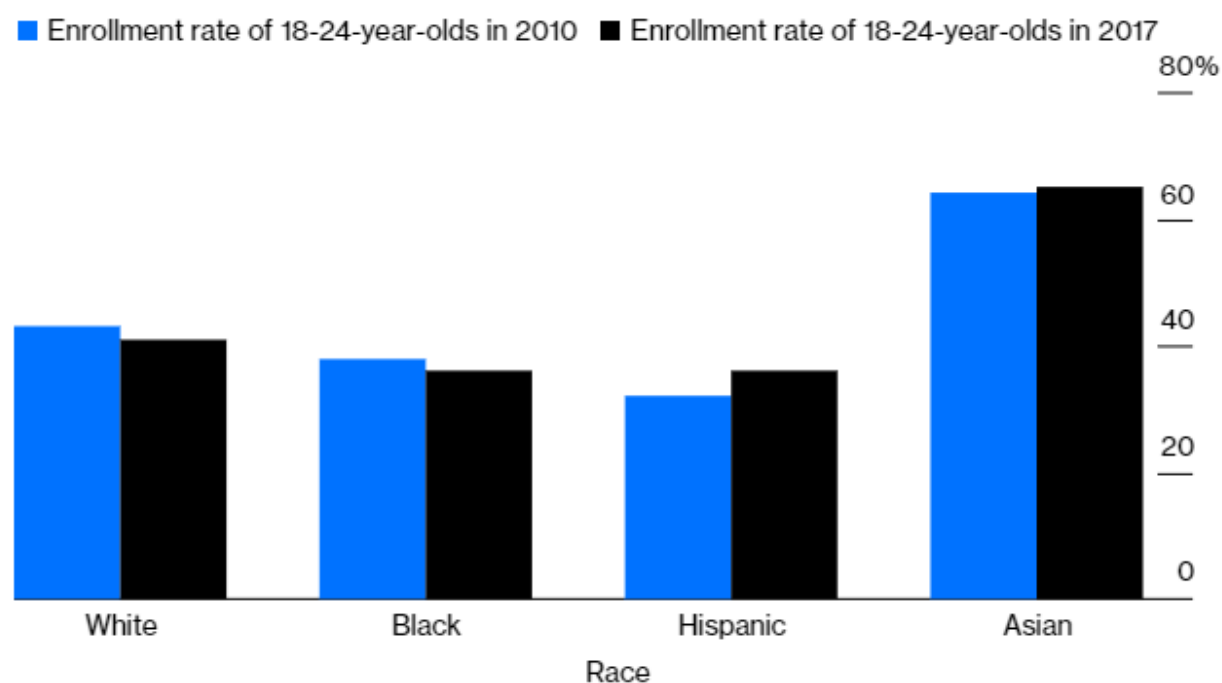
1. If the government adds a bunch of burdensome regulations or taxes, that reduces TFP.
2. If people’s education level stops increasing, that lowers TFP growth.
3. If the population gets older, that can reduce TFP (since older workers are, on average, less productive).
4. If people spend more time goofing off at work, that can reduce measured TFP (since we’re overestimating the amount of labor input).
5. If people stop moving from less productive places to more productive places (for example, if housing restrictions drive up rents and keep workers away from superstar cities), that reduces TFP.

6. If a few big companies become more dominant, that can lower TFP, either via monopoly/monopsony power, or just by reducing dynamism in the economy.
7. If demand shifts from sectors where technology is progressing rapidly (for example, manufacturing) to sectors where it's progressing slowly (e.g. services), that can reduce TFP growth, even if the rate of technological innovation in each sector remains exactly the same.

A lot of this stuff is, in fact, happening. American educational attainment is growing more slowly (after all, we can't stay in school our whole lives!). In recent years, college attendance has plateaued, even before COVID; in fact, for White and Black students it fell a bit:


A Shift in Makeup

College enrollment rates by race



Source: National Center for Education Statistics

Meanwhile, geographic mobility is way down, people are switching jobs less, and various other measures of dynamism are down. And demand has definitely shifted from physical goods to services:



As for burdensome regulation and goofing off at work, the jury is still out.

But basically we're seeing a whole lot of things happen that tend to reduce TFP growth but that have nothing to do with slowing technological progress! We can invent economically useful stuff just as brilliantly as in the past, but if the above stuff happens, TFP will still slow down.

In fact, in a recent book called “Fully Grown: Why a Stagnant Economy Is a Sign of Success”, the brilliant growth economist Dietrich Vollrath — whose [excellent blog](#) you should absolutely read — [argues that most of the slowdown](#) in TFP comes from slowing educational attainment, lower geographic mobility and economic dynamism, and and the shift from goods to services.

Now, slowing technological progress might cause decreased business dynamism (since less new tech means less fuel for new startups). But in general, most of these things are *not* caused by a slowdown in the rate of scientific discoveries or product innovations. That calls the entire [technological stagnation theory](#) into question.


And it also means that if we see a wave of inventions but slow TFP growth, it doesn't necessarily mean that those inventions didn't boost TFP! It could just mean that *other things were happening at the same time that held TFP back*. In other words, if college attendance stays flat or declines, cities keep excluding newcomers by banning construction, and consumers keep shifting to services, it could be enough to obscure most of the impact of a 2020s tech boom. But that doesn't mean the increase wasn't there! Technological innovation might be the only thing preventing TFP from slowing down *even more*, or even falling outright.

In fact, I think there's a fair chance this is what will happen in the 2020s. BUT, we should also admit that the future isn't written in stone; TFP has sped up and slowed down before, and it might do so again.

The TFP trend can change quickly

As they say, “The trend is your friend til the bend at the end.” Part of the case for future stagnation relies on looking at TFP trends and predicting that these trends will continue. But that’s not set in stone. For example, in the 1990s and early 2000s, productivity growth accelerated temporarily, to something like its old level. Here, let me borrow Dietrich Vollrath’s graph (made from Basu, Fernald and Kimball’s data):

Productivity level - utilization adjusted



Something like that red line could happen again. Maybe for another 10 or 15 years, maybe for longer. We just don’t know.

In fact, economists Nicholas Crafts and Terence Mills make exactly this point in a recent paper. They write:

Is the econometric evidence to be preferred to techno-optimism or techno-pessimism?...[F]orecast trend TFP growth based on estimated trends over the previous 20- or 25-year window exhibits considerable variation and does not show monotonic decreases...

[A]verage realised TFP growth...over 10 year intervals...has also varied substantially over time...[T]hese 10-year-ahead outturns are not well predicted by recent trend TFP growth. In particular, sharp reversals of medium-term TFP growth performance are not identified in advance. Indeed, forecasting on this basis would have missed the productivity slowdown of the 1970s, the acceleration of the mid-1990s, and the slowdown of recent years – in other words, all the major episodes during the period!

[M]edium-term TFP growth is very unpredictable. Recent performance is not a reliable guide, implying that an econometric estimate of low trend productivity growth currently does not necessarily rule out a productivity surge in the near future. The precedent of the 1990s is witness to this. [Also], a smoothed estimate of trend TFP growth has changed only slowly over time and is well above recent actual TFP growth. This suggests that pessimism about long-term prospects can easily be overdone.

In other words:

1. It's really hard to predict TFP trends.
2. If you smooth out the trends, it looks like we're due for another burst of TFP growth, though really who knows.

In other words, no matter how much evidence you marshal to show that the TFP stagnation is real, it might not indicate a *technological* slowdown, and TFP itself might have a boom soon. That doesn't prove that techno-optimism is justified for the 2020s. But it means that looking at recent TFP trends and saying "See? Stagnation is our destiny!" is definitely not warranted.

In fact, in Tyler Cowen's 2011 book *The Great Stagnation* (the most subtle and circumspect of the stagnationist books), he identifies non-technological factors as contributing to the stagnation, and he predicts that both technology and productivity growth will bounce back. Stagnationists would be well-advised to read that book!

Part 4: Science Slowdown

Low-hanging fruit and the rising cost of science

The basic idea of science stagnation is that the easiest discoveries happen first. 150 years ago, a monk sitting around playing with plants was able to discover some of the most fundamental properties of inheritance; now, biology labs are gigantic and hugely expensive marvels of technological complexity, and the NIH spends tens of billions

of dollars every year. 400 years ago we had people rolling balls down ramps to study gravity; now we study gravity with billion-dollar gravitational wave detectors that require the efforts of thousands of highly trained scientists. And so on.

In 2020, four economists — Nicholas Bloom, Charles I. Jones, John Van Reenen, and Michael Webb — published a paper quantifying this principle, and the results are deeply disturbing. Across a wide variety of fields, they found that the cost of progress has been rising steadily; more and more researchers (or “effective researchers”) are required for each incremental advance. This is exactly what a “low-hanging-fruit” model of science would predict.

Now, one objection to this type of analysis is that science doesn’t proceed within a fixed set of research fields. Bloom et al. look at specific things like Moore’s Law. But while it might take more and more research input to fit more transistors on a chip, before the 1940s there was no such thing as a transistor at all! Moore’s Law only goes back to the 1970s (some people try to generalize it and extend it further, but it really wasn’t a big driver of tech progress until relatively recently). Some fields, like agriculture, are very very old, but new ones are being born all the time. Someday someone will draw a chart showing declining research productivity in CRISPr technology, or deep learning (in fact some people are drawing such graphs even now). But those fields have been invented within our lifetime.

In reality, technological progress probably doesn’t look like a fixed set of flattening curves, but as a constantly expanding set of S-curves. We don’t just discover ways to do the same thing better; we also discover new things to do.

But Bloom et al. also look at aggregate measures — the number of effective researchers in society as a whole, versus the growth of productivity. Here’s the key graph from the paper:




FIGURE 1. AGGREGATE DATA ON GROWTH AND RESEARCH EFFORT

Even if you attribute the slowdown in TFP growth to things like slowing educational attainment, this graph still shows that we're spending more and more on research while failing to produce any noticeable acceleration in productivity — running harder and harder just to stay in place.

That's scary, because it means that eventually we'll run out of bodies to throw at scientific research, and then progress will grind to a halt. In fact, Chad Jones, one of the authors on the abovementioned paper, has already done the math on this. In a pair of papers in [1995](#) and [1999](#), he showed a simple model of what happens to growth if ideas get harder to find over time. The key result of this simple model (on page 3 of the second paper) is that constant productivity growth is *proportional to the rate of population growth*. If ideas get harder to find over time, productivity will grow more slowly than population in the long run. But even if ideas get *easier* to find as you get more of them, productivity can't grow in these models if population stagnates.

So what happens if population actually declines? In [a third paper](#), Jones shows that uninterrupted population decline causes living standards to stagnate as the planet empties out — basically the “Children of Men” scenario. With global fertility rates falling, that's definitely a worry. And since it's not clear that our available resources can even support indefinite exponential population growth, we may simply be damned-if-we-do, damned-if-we-don't; we might have no choice but to accept a long-term stagnation in living standards.


(The only salvation in this case might be *A.I. researchers*; you can imagine a Jones-type model with high labor-capital substitutability in the research production function, where it's possible to keep building machines that invent even better machines. I'm sure someone has worked out this model, and I just have to hunt around for it. **Update:** [Here is a paper](#) that does this!)

The idea of increasingly expensive research, unlike the other stagnationist arguments I addressed in earlier posts, is very hard to rebut — the theory is simple and powerful and the data is comprehensive and clear. But there are a few caveats to note here.

“More expensive” doesn't mean “slowing down”

Even if science is getting more expensive, that doesn't mean it's actually slowing down; as Bloom et al.'s graph shows, we're continuously throwing more money at research. But in [a 2018 article in *The Atlantic*](#), Patrick Collison and Michael Nielsen made a stronger argument than Bloom et al. — they argued that the rate of important scientific discoveries is actually *slowing down*. (**Update:** [Patrick clarifies](#) that they didn't want to make quite so strong an argument as I make it out to be!)

Now, both Patrick and Michael are brilliant guys (you can read [Patrick's Noahpinion interview here!](#)), and of course I agree with their overall concern. But I think they might have overstated their case a bit here. First of all, their main piece of quantitative evidence is to survey physicists about the importance of Nobel-winning discoveries from each decade:



But this is only to be expected, because *important discoveries become more important as they age*. Quantum mechanics was certainly cool stuff back in the 1920s, but it wasn't until later that things like quantum field theory built on it, or engineering applications were developed that made use of it.

Science is progressive like that; each discovery is supported by the discoveries that came before it (in Newton's words, it "stands on the shoulders of giants"). Thus, each new discovery makes the older discoveries that support it that much more important. So we'd expect to see more important discoveries in the distant past than the recent past; this is not, by itself, a sign of stagnation.

Patrick and Michael also point out that in recent years, there have been few Nobels given to research done in the 90s and 00s:

[T]he paucity of prizes since 1990 is itself suggestive. The 1990s and 2000s have the dubious distinction of being the decades over which the [Physics] Nobel Committee has most strongly preferred to skip, and instead award prizes for earlier work...

As in physics, the 1990s and 2000s are omitted [for biology and chemistry], because the Nobel Committee has strongly preferred earlier work: Fewer prizes were awarded for work done in the 1990s and 2000s than over any similar window in earlier decades.

But this is actually consistent with the narrative of an *acceleration* of Nobel-worthy discoveries in the 50s through the 80s. Remember that the Nobel Prize is rate-limited — it can only be given to a maximum of three people each year, and it has to be given while the researcher is still alive. So if there's an acceleration of Nobel-worthy discoveries, that

will create a *backlog* — a bunch of people who deserve Nobels waiting longer and longer in the queue, and a committee handing prizes to older and older researchers in order to honor them before they die. I'm not saying there *was* an acceleration of breakthrough science in the mid-20th century, but that would definitely be consistent with Nobels going to older and older discoveries.

So while I think it's possible that science is actually slowing down despite our best efforts to sustain it, I don't think we've seen evidence for that yet. Remember that "slowing down" is very different from "becoming more expensive to sustain".

Worlds enough, and time

The upshot here is that even if science is getting more expensive, we can still afford to spend more resources and sustain it for a while. There's both theory and evidence to support this proposition, both of which you can find over at [Matt Clancy's excellent blog](#).

First, some theory. In [a very readable post](#), Clancy explains [a recent theory paper](#) by Benjamin F. Jones and Larry Summers. Jones and Summers make a simple model of the economy in which R&D spending drives growth, plug recent numbers for our real economy into that model, and then ask how much it would reduce growth if we were to stop spending on R&D. The answer is: A lot.

[O]n average the return on a dollar spent on R&D is equal to the long-run average growth rate, divided by the share of GDP spent on R&D and the interest rate. With $g = 0.018$, $s = 0.025$, and $r = 0.05$, this gives us a benefits to cost ratio of 14.4. Every dollar spent on R&D gets transformed into \$14.40!

If we look more critically at the assumptions that went into generating this number, we can get different benefit-cost ratios. But the core result of Jones and Summers is not any exact number. It's that whatever number you believe is most accurate, it's much more than 1. R&D is a reliable money-printing machine: you put in a dollar, and you get back out much more than a dollar.


Clancy walks us through Jones and Summers' response to various objections, in which they add more realism to their simple model — economic growth that happens for reasons other than R&D spending, the time it takes for research discoveries to be embodied in actual technology, other costs and benefits of research besides economic growth, and so on. With any and all of these modifications, the benefits of R&D spending always exceeds the cost by a long shot.

In other words, if this model is to be believed, then spending on R&D still gets a lot of bang for the buck in our current economy — we are not near Charles Jones' imagined future dystopia in which all the low-hanging fruit of science has been picked and innovation becomes too expensive to sustain.

But a model is just a model, right? We want evidence. And in a second post, Clancy [links us to a bit of evidence](#) that R&D spending still drives progress forward, explaining the results of several papers that find that government research grants to small businesses

were very effective at creating patentable inventions in both the U.S. and the EU.

Now, patents aren't the same as dollars of GDP, and we will need more research to nail down the total economic benefits of research spending. But in terms of actual policy, we could at least spend as much of our GDP as we used to! As the Information Technology and Innovation Foundation reports, U.S. business has done its part, but federal government R&D funding has really fallen off:



In other words, business is running fast enough to stay in place in terms of R&D, but the government isn't even doing that much. That's why a big expansion in federal research funding (which until recently it looked like we were getting) is so important.

Anyway, the upshot of all this is that while the increasing cost of science is a real and significant concern, it doesn't mean it's all stagnation from here on out. We still probably have enough money to accelerate technological progress once again, if we're willing to spend it.

In Conclusion

The arguments of the techno-pessimists are not specious or ridiculous — indeed, they are very powerful. There's no denying the reality of the recent productivity slowdown, or the fact that it's getting more expensive to get scientific results (at least within specific fields). Techno-optimists have a very difficult task — predicting not only that the trends will reverse, but identifying which specific technologies will enable that reversal.

I hope I've managed to present the optimistic arguments fairly. Some of this is "optimism of the intellect" — looking at solar, batteries, and the other new technologies and predicting why they'll give productivity a bigger boost than the stuff we installed in the

2010s. And some of this is “optimism of the will” — realizing that if as a society we make a choice to commit more of our resources to pushing progress forward, we have a good chance of doing so.

To paraphrase Michael Nielsen, optimism about the future is just pessimism about the present. I look around and I’m sure we can do better than this.
