

Isolated Demands for Rigour in New Optimism

2020-12-08

I leave on vacation for a week, and you all go wild. Apparently, The Great Stagnation ended while I was away, and we are now celebrating a brave new era of progress.

Graciously compiled by Caleb Watney, we have:

- Noah Smith, Techno-optimism for the 2020s
- Matthew Yglesias, Some optimism about America's Covid response
- Tyler Cowen, Is the Great Stagnation over?
- Caleb Watney, Cracks in the Great Stagnation

This recent wave of anti-stagnation writing consists largely of anecdotes, supported by scant evidence, and without any sign of serious thought.

The problem is not even that the ideas are wrong.

The problem is the blatantly imbalanced and isolated demands for rigour.

Three years ago, to be convinced that there was actually a Great Stagnation, we required:

- Are Ideas Getting Harder to Find, a 54 page paper published in American Economic Review, co-authored by two Stanford Professors, an MIT professor and a Stanford PhD student.
- Science Is Getting Less Bang for Its Buck, an Atlantic article co-authored by Patrick Collison and Michael Nielsen.
- Tyler Cowen and Ben Southwood's 41 page monograph, filled with original data analysis, detailing the slowing of science across numerous fields.

And of course, The Great Stagnation itself, a book length treatment of the question.

These were not luke-warm proposals. Bloom et al concludes **“Our robust finding is that research productivity is falling sharply everywhere we look.”** Collison and Nielsen write: “the evidence is that science has slowed enormously per dollar or hour spent”. And from Cowen and Southwood:

To sum up the basic conclusions of this paper, **there is good and also wide-ranging evidence that the rate of scientific progress has indeed slowed down**, [in] the disparate and partially independent areas of productivity growth, total factor productivity, GDP growth, patent measures, researcher productivity, crop yields, life expectancy, and Moore's Law we have found support for this claim.

In contrast, each of the new anti-stagnation posts consists largely of one-off innovations that the authors happen to think are cool. There is no evidence that they are indicative of more generalized trends. Noah Smith leads with a meme comparing Juicero to SpaceX Starship, Caleb throws out a heartwarming video of a self-driving Tesla.

I don't think it's controversial to say that the pro-stagnation evidence was of dramatically higher caliber.

But today, all it takes is some guy tweeting out a few bullet points, and everyone loses their minds. Seriously, how credulous are you?

Has the Evidence Shifted?

If you actually care about overturning the stagnation hypothesis, start with the papers that proved it in the first place.

Is science accelerating once again? Are we in a new era of progress? If so, it should show up in the data.

In 2017, Bloom et al concluded that ideas were getting harder to find on the basis of R&D data from semiconductor companies. By inferring headcount from R&D spending, they produced this chart:





FIGURE 4. DATA ON MOORE'S LAW

What's changed since in the last 3 years?

I don't have access to Bloom's data provider, but I was able to grab Intel's recent R&D spending from their 10-K filings, and append it onto Bloom's historical data:

Intel Nominal R&D Spend (\$MM)




It does actually look like costs might be slowing! Hurray for progress, and hurray for the end of stagnation!

But wait a minute, Intel is no longer the most Moore's Law-relevant company. Their 7nm process was delayed to 2022, and they no longer lead the pack.

Instead, TSMC is now one of only two fabs (including Samsung) able to keep up with Moore's Law. (For what it's worth, they also manufactured the Apple M1 chip.) This is their R&D data from Bloom, with the last few years added.

TSMC Nominal R&D Spend (\$MM)




Data [here](#), collected from [TSMC SEC Filings](#).

Costs continue to rise exponentially, with no clear break in trend since 2017. If you believed the Bloom et al paper when it was published three years ago, you have no legitimate reason to turn your back on it now.

What about the other fields Cowen/Southwood cite as proof of stagnation?

Here's their data on life expectancy up to 2016:


Figure 2 Life expectancy at birth in the US, 1800-2016



As you can see, there's a plateau starting around the early 2000s, with expectancy actually dropping off in the most recent data. When you all saw this data in 2019, you seemed pretty darn convinced that science was slowing down. Upon reading the monograph in its entirety, Scott Alexander wrote:

I had previously argued technological progress wasn't slowing down; based on the work of Tyler Cowen and Ben Southwood I now think it is; my previous position was mistaken.

Okay, so the data up until 2016 was convincing. What's happened since then? Here's life expectancy again, extended to include data up to 2020:



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.

Is the New Evidence Exceptionally Compelling?

So fine, none of those factors have changed. But surely these new discoveries still constitute categorical shifts in our trajectory?

Let's take a closer look. Aggregated across the aforementioned posts plus more from [Tyler Cowen](#) and [Eli Dourado](#), here is some of the evidence presented in favor of progress (full list in the appendix):


1. Affordable Solar Power
2. Crypto Going Mainstream + Ethereum 2.0
3. China Announces Quantum Supremacy

1. Affordable Solar Power

Here's the cost per watt for solar power:

Solar PV module prices

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



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

OurWorldInData.org/energy • CC BY

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?

Noah Smith writes:

...the price drops for solar and wind over the last decade have just been nothing short of revolutionary.

Caleb Watney similarly writes:

the almost ho-hum daily progress in solar, wind, and battery technology where prices have fallen 90, 70, and 87 percent over the last ten years

So both are citing work over the last 10 years, rather than more recent progress. What's going on here?

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?

More likely, solar power has been making great strides on a pretty consistent basis for decades, and the only recent break is in how high-status it is to say that out loud.

2. Mainstream Crypto / Ethereum 2.0


Again, a big deal, and certainly neat to have crypto adopted by PayPal.

Except that it hasn't actually happened yet. And also, crypto was already adopted by Stripe, before being subsequently shut down.


And again, I have to ask, is this a bigger deal than when Bitcoin first launched in 2009? Is it a bigger deal than Ethereum 1.0 first going live in 2015?

Here's some actual data on adoption and usage:

Ethereum Transactions Per Day



Bitcoin Transactions Per Day



Again, please explain to me how 2020 was the year crypto went mainstream. Please explain why you were unconvinced by crypto as a sign of progress in 2017, but have since changed your mind.

3. Quantum Supremacy

- 2011: D-Wave sells a quantum computer to Lockheed Martin, you all say there is a great stagnation
- 2019: Google achieves quantum supremacy, you all say there is a great stagnation

- 2020: China achieves quantum supremacy (but this time with photonics), you all lose your minds

Here are the Scott Aaronson posts for [Google](#) and for [China](#). He explains that the latter:

...represents the first demonstration that quantum supremacy is possible via photonics. Finally, as the authors point out, the new experiment has one big technical advantage compared to Google's: namely, many more possible output states (~1030 of them, rather than a mere ~9 quadrillion). This makes it infeasible to calculate the whole probability distribution over outputs and store it on a gigantic hard disk (after which one could easily generate as many samples as one wanted), which is what IBM proposed doing in its response to Google's announcement.

So yes, it is a new achievement. Is it a demonstration of progress in a way that the Google achievement was not? I don't understand enough about quantum computing to say, but critically, neither do any of you.

My point isn't to celebrate ignorance. It's to say that the 2019 announcement was clearly a categorical leap forward. Somehow, that was not enough to convince you all of scientific progress.

This year, we have another advance. Because there is not a clear index here, and because the 2019 achievement was so tremendous, it's difficult to say whether or not this is a break in trend.

Again, unless you have a strong belief otherwise, you should either take both achievements as a sign of progress, or neither. You cannot merely pick and choose what to celebrate based on what is trendy rhetoric.

Conclusion

I say this not because I am "against progress", but because I am very much in your corner of the internet. I also get all of my news from Marginal Revolution, and consider myself part of the broader Progress Studies community.

And so speaking as an insider, I am saying that we have to do better.

The recent trend of rhetoric against stagnation is not founded in evidence or driven by data. It is pure mood affiliation. You bought into the stagnation hypothesis when it was hip and contrarian, and now buy into the optimism hypothesis to be even more hip and counter-contrarian. At no point did you stop to look at the data or actually think for yourself.

To be clear, none of this is to say that The Great Stagnation is not over! Maybe yes, maybe no, maybe the whole thing was an illusion and reality is driven by the gods of straight lines.

The same month Cowen and Southwood published their monograph on stagnation, Alexey Guzey wrote:

science is not slowing down... I think that the perception of stagnation in science – and in biology specifically – is basically fake news, driven by technological hedonic treadmill and nostalgia. We rapidly adapt to technological advances – however big they are – and we always idealize the past – however terrible it was.

If anyone now has the moral authority to speak on optimism (or to identify as a contrarian with a straight face), it is him.

And if you sincerely believe that we are in a new era of progress, then argue for it rigorously! Show it in the data. Revisit the papers that were so convincing to you a year ago, and go refute them directly.

Maybe then I will be happy to celebrate alongside you.

TLDR

Go read this [SMBC comic](#). Then read [this SlateStarCodex post](#). Then go look at yourself in the mirror.

Appendix: Notes on Herd Mentality and Mood Affiliation

As much as I hesitate to make essentialist remarks, I do believe that there are fundamentally two types of people. Those who see a bunch of people agree, and think “wow, everyone believes this, it must be correct!”

And then those who think “hmm, suspicious”.

It’s not that I’m positing any kind of coordinated conspiracy, it’s that I don’t even have to. Because humans are basically apes and prone to shallow mimicry, it only takes one very prominent ape to have an opinion, and everyone else will rush to share it.

More specifically, additional opinion pieces don’t qualify as additional evidence for your cause unless they actually make different points than each other. All four of those articles cite the Moderna vaccine as evidence. That’s fair, it is a strong piece of evidence. But as a reader, you get to count it exactly once.

To make matters worse, many of the pieces cite each other. Noah reads Caleb, who reads Tyler. If you trace the intellectual lineage, it quickly becomes clear that approximately one person has any original ideas, and everyone else is just piling on.

The Wisdom of Crowds only functions when each individual is capable of thinking and acting independently. Absent this vital condition, it is just madness.

Appendix: All Innovations Cited in Favor of Progress

Aggregated across all sources listed above, here are all the innovations:

- mRNA Vaccine
- Apple M1 Chip
- SpaceX Launch / SpaceX Starship (delayed until Monday)

- GPT-3 / AI
- Electric Cars
- Mainstream Crypto / Ethereum 2.0
- Operation Warp Speed
- Affordable Solar Power / Green Energy
- The Eggplant
- Remote Work
- V-Shaped Recovery
- Tons of cool companies IPO'ing and tons more getting started,
- DeepMind Protein Folding
- Lab-Grown Meat Lab Grown Meat Approval
- Sight Restored in Mice

I've done a few already, showing that the rest are not an obvious departure from existing trends is left as an exercise to the reader.

To be clear, the question is not “is this innovation very cool”, but rather “does this innovation depart from the previous decade's trend of progress”.

For example:

- The mouse study is cool, but is this a bigger deal than the 2012 discovery of CRISPR-Cas9 programming?
- GPT-3 is very cool, but we've been on a sharp trajectory of progress in ML since 2012. GANs have seen enormous progress every year since 2014, as have many other tasks.
- The Starship hop is very impressive. But is this a bigger leap forward than in 2008 when SpaceX became the first private company to ever launch, orbit and recover a spacecraft? Is it bigger than in 2015 when they achieved the first vertical landing, or in 2017 when they achieved the first vertical landing of an orbital rocket?
- Electric cars have seen substantial progress with more competition from mainstream automakers, but surely the biggest breakthroughs were the General Motors EV1 in 1996, the Tesla Roadster in 2008 and perhaps the Tesla Model 3 in 2017? Here's some actual data, if you care at all about that kind of thing. It demonstrates steady progress, with no clear inflection point or recent change of trajectory.

Global electric car stock, 2010-2019

Open 

