"So the real scandal is: Why did anyone ever listen to this guy?"

Posted on May 8, 2020 9:27 AM by Andrew

John Fund writes:

> *[Imperial College epidemiologist Neil] Ferguson was behind the disputed research that sparked the mass culling of eleven million sheep and cattle during the 2001 outbreak of foot-and-mouth disease. He also predicted that up to 150,000 people could die. There were fewer than 200 deaths. . . .*
>
> *In 2002, Ferguson predicted that up to 50,000 people would likely die from exposure to BSE (mad cow disease) in beef. In the U.K., there were only 177 deaths from BSE.*
>
> *In 2005, Ferguson predicted that up to 150 million people could be killed from bird flu. In the end, only 282 people died worldwide from the disease between 2003 and 2009.*
>
> *In 2009, a government estimate, based on Ferguson's advice, said a "reasonable worst-case scenario" was that the swine flu would lead to 65,000 British deaths. In the end, swine flu killed 457 people in the U.K.*
>
> *Last March, Ferguson admitted that his Imperial College model of the COVID-19 disease was based on undocumented, 13-year-old computer code that was intended to be used for a feared influenza pandemic, rather than a coronavirus. Ferguson declined to release his original code so other scientists could check his results. He only released a heavily revised set of code last week, after a six-week delay.*
>
> *So the real scandal is: Why did anyone ever listen to this guy?*

I don't know. It's a good question. When Ferguson was in the news a few months ago, why wasn't there more discussion of his atrocious track record? Or was his track record not so bad? A google search turned up this op-ed by Bob Ward referring to Ferguson's conclusions as "evidence that Britain's political-media complex finds too difficult to accept." Regarding the foot-and-mouth-disease thing, Ward writes, "Ferguson received an OBE in recognition for his important role in the crisis, or that he was afterwards elected a fellow of the prestigious Academy of Medical Sciences." Those sorts of awards don't cut much ice with me—they remind me too much of the U.S. National Academy of Sciences—but maybe there's more of the story I haven't heard.

I guess I'd have to see the exact quotes that are being referred to in the paragraphs excerpted above. For example, what did Ferguson exactly say when he "predicted that up to 150,000 people could die" of foot-and-mouth disease. Did he say, "I expect it will be under 200 deaths if we cull the herds, but otherwise it could be up to 2000 or more, and worst case it could even be as high as 150,000?" Or did he flat out say, "150,000, baby! Buy your gravestone now while supplies last."? I wanna see the quotes.

But, if Ferguson really did have a series of previous errors, then, yeah, Why did anyone ever listen to this guy?

In the above-linked article, Fund seems to be asking the question rhetorically.

But it's a good question, so let's try to answer it. Here are a few possibilities:

1. Ferguson didn't really make all those errors; if you look at his actual statements, he was sane and reasonable.

Could be. I can't evaluate this one based on the information available to me right now, so let's move on.

[Indeed, there seems to be some truth to this explanation; see P.S. below.]

2. Nobody realized Ferguson had made all those errors. That's true of me—I'd never heard of the guy before all this coronavirus news.

We may be coming to a real explanation here. If a researcher has success, you can find evidence of it—you'll see lots of citations, a prestigious position, etc. But if a researcher makes mistakes, it's more of a secret. Google the name and you'll find some criticism, but it's hard to know what to make of it. Online criticism doesn't seem like hard evidence. Even published papers criticizing published work typically don't have the impact of the original publications.

3. Ferguson played a role in the system. He told people what they wanted to hear—or, at least, what some people wanted to hear. Maybe he played the role of professional doomsayer.

There must be something to this. You might say: Sure, but if they wanted a doomsayer, why not find someone who hadn't made all those bad predictions? But that misses the point. If someone's job is to play a role, to speak from the script no matter what the data say, then doing bad work is a kind of positive qualification, in that it demonstrates one's willingness to play that role.

But this only takes us part of the way there. OK, so Ferguson played a role. But why would the government want him to play that role. If you buy the argument of Fund (the author of the above-quoted article), the shutdowns were a mistake, destructive economically and unnecessary from the standpoint of public health. For the government to follow such advice—presumably, someone must have been convinced of Ferguson's argument from a policy perspective. So that brings us back to points 1 and 2 above.

4. A reputational incumbency effect. Once someone is considered an expert, they stay an expert, absent unusual circumstances. Consider Dr. Oz, who's an expert because people consider him an expert.

5. Low standards. We've talked about this before. Lots of tenured and accoladed professors at top universities do bad work. I'm not just talking about scandals such as pizzagate or that ESP paper or epic embarrassments such as himmicanes; I'm talking more about everyday mediocrity: bestselling books or papers in top journals that are constructed out of weak evidence. See for example here, here, and here.

The point is, what it takes to be a celebrated academic is to have some successes. You're defined by the best thing you did, not the worst.

And maybe that's a good thing. After all, lots of people can do bad work: doing bad work doesn't make you special. I proved a false theorem once! But doing good work, that's something. Now, some of these celebrity academics have never done any wonderful work, at least as far as I can tell. But they're benefiting from the general principle.

On the other hand, if the goal is policy advice, maybe it's better to judge people by their worst. I'm not sure.

Not that we're any better here in the U.S., where these academics have had influence in government.

Taking the long view, organizations continue to get staffed with knaves and fools. Eternal vigilance etc. Screaming at people in the press isn't a full solution, but it's a start.

P.S. There seems to some truth to explanation 1 above, "Ferguson didn't really make all those errors; if you look at his actual statements, he was sane and reasonable." From Tom in comments:

> *Mad Cow paper:*
> *https://www.ncbi.nlm.nih.gov/pubmed/11786878*
>
> *From abstract:*
> *"Extending the analysis to consider absolute risk, we estimate the 95% confidence interval for future vCJD mortality to be 50 to 50,000 human deaths considering exposure to bovine BSE alone, with the upper bound increasing to 150,000 once we include exposure from the worst-case ovine BSE scenario examined."*
>
> *Consistent with the "up to 50,000" quote but the quote fails to mention the lower bound.*

See also Vidur's comment which discusses some of the other forecasts.

This entry was posted in Zombies by Andrew. Bookmark the permalink.

---

652 THOUGHTS ON ""SO THE REAL SCANDAL IS: WHY DID ANYONE EVER LISTEN TO THIS GUY?""

James Annan on **May 8, 2020 9:32 AM at 9:32 am** said:

Ferguson has made some errors (not calibrating his model is one of the largest) but most of this is bogus. Saying "if you don't act, X will (or at least may) happen" is not a wrong prediction if people take action to avoid X, and then X fails to happen.

Reply ↓

Carlos Ungil on **May 8, 2020 9:44 AM at 9:44 am** said:

What do you mean by "calibrating his model"?

Reply ↓

Andrew [not Gelman] on May 8, 2020 4:14 PM at 4:14 pm said:

I think he means that Ferguson et al assumed a doubling time of 5 days at a time when a fit to available data would have shown it to be less than 3 days. This meant that they may have under-estimated R0, which in turn may explain why the current UK deaths exceed Ferguson's forecasts for the lockdown scenario.

Reply ↓

Carlos Ungil on May 9, 2020 3:44 PM at 3:44 pm said:

Thanks. I've just noticed that I could have find out following the link in his name.

Reply ↓

dhogaza on May 9, 2020 4:09 PM at 4:09 pm said:

It looks like the team working on the C++ version are calibrating the model. I have no idea how it's being used, though.

"…. Right now we are in the emergency stage where model runs are repeated multiple times in the week with calibrations to the latest data as it comes."

https://github.com/mrc-ide/covid-sim/issues/194

Reply ↓

Marilyn Shepherd on May 22, 2020 6:32 AM at 6:32 am said:

Doing some other research it didn't take long to discover he based his "model" on Spanish flu, a different virus from over 100 yrs ago and thought this one might be the same. Coronas have been around forever without a single nation ever shutting down.

Reply ↓

V on November 5, 2020 9:18 AM at 9:18 am said:

NB. MERS and SARS were both new corona viruses. This one is a new corona virus. Virologists are the best people to explain and identify viruses. Beware of misleading inaccurate sweeping generalisations.

Reply ↓

Daniel Connelly on June 15, 2020 11:41 AM at 11:41 am said:

You need to be careful about calibrating the model to doubling time, since early doubling time can be dominated by a subset of the population, while herd immunity may be determined by a broader component of the population. For example, those living in dense urban areas may dominate the early increase, but assuming this same R for herd immunity would lead to an overestimate of net infections. So if the calibration was conservative on R, that may have yielded a better final result.

That said, I've not seen anything indicating the ultimate death numbers in the absence of interventions would be substantially (within a factor of 2) off from the study's estimates. And a key conclusion of the paper is that relaxing interventions prematurely (before a vaccine, or aggressive test & trace capability) would result in a return of exponential growth.

Reply ↓

Twain on May 8, 2020 9:51 AM at 9:51 am said:

James Annan,

> Saying "if you don't act, X will (or at least may) happen" is not a wrong prediction if people take action to avoid X, and then X fails to happen.

Are you claiming that "If we did X (shutdowns) to avoid Y (deaths from COVID-19) and then saw Y not occur, X must be the cause"?

If yes, that is the Post Hoc Fallacy: "Since event Y followed event X, event Y must have been caused by event X."

Reply ↓

Zhou Fang on May 8, 2020 2:22 PM at 2:22 pm said:

He's not saying X must be the cause, he's saying you cannot invalidate a conditional prediction if the condition never takes place.

Reply ↓

Linch Zhang on May 8, 2020 9:59 AM at 9:59 am said:

No, he's saying that if somebody makes a claim "If X, then Y" and neither X nor Y happened, this can't be much evidence for the causal chain being wrong unless you have fairly sophisticated machinery.

Reply ↓

Twain on May 8, 2020 10:26 AM at 10:26 am said:

Linch,

I don't follow.

Reply ↓

Carlos Ungil on May 8, 2020 10:46 AM at 10:46 am said:

This was the prediction:

"In the (unlikely) absence of any control measures or spontaneous changes in individual behaviour, we would expect a peak in mortality (daily deaths) to occur after approximately 3 months [...] given an estimated R0 of 2.4, we predict 81% of the GB and US populations would be infected over the course of the epidemic. [...] In total, in an unmitigated epidemic, we would predict approximately 510,000 deaths in GB and 2.2 million in the US, not accounting for the potential negative effects of health systems being overwhelmed on mortality."

We cannot say that the prediction was wrong just because, after a few weeks and with very strong mitigation measures in place, death have been only 31,000 and 77,000.

Reply ↓

Joshua on May 8, 2020 10:51 AM at 10:51 am said:

Carlos –

Not wanting to be overly semantic – but I think the distinction between prediction and projection is important here.

Reply ↓

Carlos Ungil on May 8, 2020 11:04 AM at 11:04 am said:

I don't know what that means but the choice of word is theirs, not mine. I just relayed the "prediction".

Reply ↓

Joshua on May 8, 2020 11:11 AM at 11:11 am said:

In climate change, a result of a model often implies a projection – say that if emissions are at X level going forward, temps will rise within a range with Y being at the top of that range.

People with an agenda will then say that "The model predicted that temps will reach Y, and temps didn't reach Y, therefore the modelers were trying to scare us to destroy capitalism, and btw, modeling is bad."

The not only ignore that emissions were a X-n, they also ignore the lower bound of the projection.

I think that a big part of the misundertanding/misuse of the models lies in the distinction between a "prediction" and a <conditional) "projection."

Lowell on September 19, 2020 1:47 PM at 1:47 pm said:

77,000 dead in the US from this plandemic?

CDC reported 2018 influenza deaths at 79,000
CDC reported 2018 pneumonia deaths at 40,000
Now CDC says just less than 11,000 are attributable to COVID-19 alone.

Comorbidities (https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm)
Table 3 shows the types of health conditions and contributing causes mentioned in conjunction with deaths involving coronavirus disease 2019 (COVID-19). For 6% of the deaths, COVID-19 was the only cause mentioned. For deaths with conditions or causes in addition to COVID-19, on average, there were 2.6 additional conditions or causes per death. The number of deaths with each condition or cause is shown for all deaths and by age groups. For data on comorbidities,

Some categorization issues are clouding the real data.

Saying as many as NONE to 50,000 would make a really unpopular fortune teller, but profitable in other beliefs.

Reply ↓

Ian Kenny on June 5, 2021 10:20 AM at 10:20 am said:

Why cannot you say that the prediction was wrong, personally I think as he was out by 479,000 and 2.123 million it was moronic

Reply ↓

Daniel Lakeland on May 8, 2020 10:49 AM at 10:49 am said:

If you shoot yourself in the leg, you will bleed to death. Later… you don't shoot yourself in the leg… and you do not bleed to death… headlines read "Twain does not die of exsanguination proving Daniel has no idea what he's talking about"

Reply ↓

Twain on May 8, 2020 11:05 AM at 11:05 am said:

Daniel,

> If you shoot yourself in the leg, you will bleed to death. Later… you don't shoot yourself in the leg… and you do not bleed to death… headlines read "Twain does not die of exsanguination proving Daniel has no idea what he's talking about"

This is a strawman argument that does not represent the argument I am presenting.

My argument is this: One cannot say "We implemented shutdowns (X) and cases/deaths from COVID-19 (Y) decreased. Therefore X is the sole cause of Y."

This ignores the numerous factors exclusive the the shutdowns — improved hygiene (before and during shutdowns), masks, voluntary distancing (prior to shutdown), immunity, weather, etc. — that could be causing Y in addition to X. Sure, X could have and affect — but say it is the sole cause is a logical fallacy in absence of decisive data, which we do not have yet.

Reply ↓

Daniel Lakeland on May 8, 2020 11:24 AM at 11:24 am said:

You're still missing the point… if I predict that shooting yourself in the leg will cause you to bleed to death… and later you don't bleed to death… a sensationalist/politicized journalism article could come out saying "Daniel predicts Twain will bleed to death… but he doesn't! so Daniel has no idea what he's talking about"

but I didn't predict "Twain will bleed to death" I predicted "IF YOU SHOOT YOURSELF then you will bleed to death". Since you DIDN'T SHOOT YOURSELF my prediction has yet to be tested.

This seems to be a consistent thread here, that frequently a modeler creates a *conditional prediction*… "under conditions X then Y will happen" and it's taken as an *unconditional* prediction of "Y will happen, panic!" later when Y doesn't happen, it's taken as evidence that modeler has no idea what he's talking about (usually for political reasons, to attempt to discredit them).

The only way you can disprove the goodness of a conditional probabilistic prediction is to see the conditions X come true, and then see that actual outcome Y is fairly far outside the high probability range of Y predicted by the model.

So for example in the prediction quoted by Carlos above: https://statmodeling.stat.columbia.edu/2020/05/08/so-the-real-scandal-is-why-did-anyone-ever-listen-to-this-guy/#comment-1331888

Since we did not see "In the (unlikely) absence of any control measures or spontaneous changes in individual behaviour" happen, it is now impossible to determine whether that was a good prediction by looking at the observations. All we can do is look at whether the mechanisms used in the prediction, or the software code to calculate it or anything like that might have had bugs or bad assumptions. The actual outcomes are in a counterfactual world where the conditions simply didn't hold and are therefore irrelevant to the accuracy of the prediction.

Reply ↓

CROM on August 7, 2020 8:25 AM at 8:25 am said:

You conveniently forget that the specific modeler overshot reality in several past instances too.

“ In 2005, Ferguson said that up to 200 million people could be killed from bird flu. He told the Guardian that ‘around 40 million people died in 1918 Spanish flu outbreak… There are six times more people on the planet now so you could scale it up to around 200 million people probably.’ In the end, only 282 people died worldwide from the disease between 2003 and 2009.”

He was responsible for the killing of some millions of animals for foot and mouth disease

“ In 2001 the Imperial team produced modelling on foot and mouth disease that suggested that animals in neighbouring farms should be culled, even if there was no evidence of infection. This influenced government policy and led to the total culling of more than six million cattle, sheep and pigs – with a cost to the UK economy estimated at £10 billion.
It has been claimed by experts such as Michael Thrusfield, professor of veterinary epidemiology at Edinburgh University, that Ferguson’s modelling on foot and mouth was ‘severely flawed’ and made a ‘serious error’ by ‘ignoring the species composition of farms,’ and the fact that the disease spread faster between different species. Does Ferguson acknowledge that his modelling in 2001 was flawed and if so, has he taken steps to avoid future mistakes?”

, some ludicrous numbers from mad cows disease

“ In 2002, Ferguson predicted that between 50 and 50,000 people would likely die from exposure to BSE (mad cow disease) in beef. He also predicted that number could rise to 150,000 if there was a sheep epidemic as well. In the UK, there have only been 177 deaths from BSE.”

What more do you need to admit that models are incredibly susceptible to GIGO and modelers are incredibly unwise/ignorant to prevent GIGO?

Twain on **May 8, 2020 11:33 AM at 11:33 am** said:

Daniel,

Thank you for clarifying. Apologies, I misunderstood your point.

I agree with everything you stated.

Reply ↓

Dyke on **March 27, 2021 12:49 PM at 12:49 pm** said:

What you are trying to say is one of the basic rules in logic – only consequent is certain (result), antecedent (cause) can be anything.

Reply ↓

Twain on **May 8, 2020 10:53 AM at 10:53 am** said:

Carlos,

Thank you for clarifying.

> We cannot say that the prediction was wrong just because, after a few weeks and with very strong mitigation measures in place, death have been only 31,000 and 77,000.

Fair. We can't say it is right, either, just yet. So it is TBD.

I will add that Ferguson's above prediction (if my memory is correct) assumes a uniform IFR applied to infected individuals — which as many have discussed here at length, is not a reasonable assumption because SARS-CoV-2 has an IFR that skews heavily with age and comorbitities.

Reply ↓

jim on May 8, 2020 11:10 AM at 11:10 am said:

What we can say is that the forecast/prediction is framed responsibly with many important constraints and caveats, even giving the R0 used to make it.

For my money, if the quote is accurate, then this is a reasonable and responsible forecast. One can quibble with the actual numbers or methods, but the context is clear and the forecast is framed responsibly. He should be credited for that.

Reply ↓

Twain on May 8, 2020 11:16 AM at 11:16 am said:

Jim,

I disagree.

By not making clear his calculations assume a uniform IFR to calculate deaths for a disease where IFR varies tremendously (0.005% to 20%…) depending on age, comorbidity, etc., Ferguson misleads readers.

Not making key assumptions clear is not responsible.

Reply ↓

Daniel Lakeland on May 8, 2020 11:32 AM at 11:32 am said:

I'm not clear actually on whether he uses a uniform IFR, can you provide a source for this assertion? but you should also remember that those predictions were made at a much earlier stage in the pandemic when information available was much more limited.

I agree with jim. For the most part, everything here looks quite responsible. Even using a single IFR is perfectly fine when calculating results for an entire population. the key is just whether the population averaged IFR is close to the chosen IFR.

Twain on May 8, 2020 11:51 AM at 11:51 am said:

Daniel,

I'm assuming he is using a uniform IFR because that is what the Imperial College London Model uses per their last report (if I'm wrong, please skewer me!).

Can you explain how using a population-averaged IFR is reasonable?

To me, it seems unreasonable because IFR varies by orders-of-magnitude and skews to a specific population — so is an "average" IFR even possible? If the variance was

smaller and skewing was less, then I'd agree — but that isn't the case.

Is there something I am missing?

dhogaza on May 8, 2020 12:07 PM at 12:07 pm said:

Twain ...

"By not making clear his calculations assume a uniform IFR to calculate deaths for a disease where IFR varies tremendously (0.005% to 20%...) depending on age, comorbidity, etc., Ferguson misleads readers."

Are we being subjected to a repeat of the "mostly old people would die so just isolate them and let everyone else lead their lives normally" argument here?

Also, please do remember that death isn't the only consequence here.

Carlos Ungil on May 8, 2020 12:26 PM at 12:26 pm said:

The do not assume a uniform IFR. See page 5 in https://spiral.imperial.ac.uk:8443/bitstream/10044/1/77482/14/2020-03-16-COVID19-Report-9.pdf

dhogaza on May 8, 2020 12:26 PM at 12:26 pm said:

> Are we being subjected to a repeat of the "mostly old people would die so just isolate them and let everyone else lead their lives normally" argument here?

No.

I'm stating that Ferguson did not communicate a key underlying assumption of his model. There is a big difference, IMO, between using a uniform IFR versus an age-stratified IFR since it varies so much; the reader should know that information. He should have prefaced with "Assuming a uniform IFR for the entire population, I conditionally project..."

> Also, please do remember that death isn't the only consequence here.

No sure what you imply here.

Twain on May 8, 2020 12:28 PM at 12:28 pm said:

Ignore the above post; my mistake.

> Are we being subjected to a repeat of the "mostly old people would die so just isolate them and let everyone else lead their lives normally" argument here?

No.

I'm stating that Ferguson did not communicate a key underlying assumption of his model. There is a big difference, IMO, between using a uniform IFR versus an age-stratified IFR since it varies so much; the reader should know that information. He should have prefaced with "Assuming a uniform IFR for the entire population, I conditionally project..."

> Also, please do remember that death isn't the only consequence here.

No sure what you imply here.

Daniel Lakeland on **May 8, 2020 12:49 PM at 12:49 pm** said:

Thanks Carlos!

So now once again we conclude.... totally responsible messaging from the researcher, appropriate assumptions, irresponsibly spun or distorted in media and popular culture for political purposes.

Daniel Lakeland on **May 8, 2020 1:01 PM at 1:01 pm** said:

> Can you explain how using a population-averaged IFR is reasonable?

So, as carlos points out *no* they didn't use an average IFR.

But since IFR is a number between 0 and 1, the average over the actual population exists. if you are modeling the entire population, then you can use whatever that average is...

Suppose you have k age groups... and IFR_i is the IFR for each age group, *and* you assume that all age groups are equally susceptible and at risk to infection. Suppose Nt is the total population. N_i is the number of people in group i, and P_i is the proportion of group i that gets infected. Then the total number of people who die is:

dead = sum(IFR_i * P_i * N_i)

dead/Nt = sum(IFR_i * P_i * N_i/Nt)

Call P_i * N_i/Nt the probability of a random person being infected. pinf_i

dead/Nt = sum(IFR_i * pinf_i)

now multiply by Nt

dead = sum(IFR_i * pinf_i) * Nt

call sum(IFR_i*pinf_i) as IFR_avg.

dead = IFR_avg * Nt

so provided you calculate it appropriately, IFR_avg produces the same results as stratifying explicitly.

Twain on **May 8, 2020 1:05 PM at 1:05 pm** said:

Carlos,

Yes, I'm going to read something before making a claim.

They do assume and apply a uniform IFR; see pg 22 of Report 20.

Specifically, IFR' ~ IFR * N(1,0.1). They then apply this IFR' to calculate deaths: d = IFR' * summation(c * pi,t), where c is number of cases and pi is incubation period. PDF N(1,0.1) is to "incorporate the uncertainty inherent in this estimate we allow the IFR for every region to have additional noise around the mean."

The above, translated in practical terms, means "anyone infected with SARS-CoV-2 has a probability of dying with mean 0.9% with sigma = +/-0.1%.

So they are therefore claiming that the population in the US, who is age 0-49, has that probability of dying. Yet for any subgroup of this range of ages, per their Report 9, does not exceed 0.15% (and does not exceed 0.05% for age 0-29).

That makes no physical sense based on the data. The majority of children, teenagers, healthy adults (or those with well-managed conditions), contract SARS-CoV-2 and experience no to mild symptoms. To say the have almost a 1% chance of dying is absurd.

Twain on May 8, 2020 1:24 PM at 1:24 pm said:

Daniel,

Thank you for the detailed explanation. I appreciate it and I was mistaken with my earlier comments, sorry about that.

One minor question: IFR_avg is a direct function of accurate IFR_i and P_i and therefore sensitive to these values? Therefore, IFR_avg is a direct function of testing and accurate count of infections?

Carlos Ungil on May 8, 2020 1:36 PM at 1:36 pm said:

Twain, they have different models for different things. For calculating those 500'000 deaths under dicussion they were not using a uniform IFR in the epidemiological model simulation. I think that's pretty clear.

Anyway, the fact that you can point to the place in the more recent report on Italian regions where they discuss their assumptions suggests they are not hiding them, right? This is a completely different problem, where they need to infer infection rates from observed deaths and a simpler model may be adequate for that purpose. That same average IFR can later be used for the predictive (or projective, whatever) part if the population infected going forward is similar to the population infected in the past. Even if you were using age-specific IFRs and estimates of infection reates you would still need to aggregate them in the end. Considering that the infected populations in the first and second waves are similar is probably not the wildest implicit or explicit assumption in the model, to be fair.

Twain on May 8, 2020 1:54 PM at 1:54 pm said:

Carlos,

See my above comment in response to Daniel's explanation/derivation.

Stephen M St Onge on May 11, 2020 6:00 PM at 6:00 pm said:

I think you are being unreasonable here. If I ask "How many Los Angeles residents will die in the case of a magnitude 8.0 earthquake with epicenter downtown?", the answer I am looking for is a single number. In this scenario, I'm asking how many people die in the city limits, how many in the various suburbs of LA, how the San Fernando Valley will differ from the region south of the Santa Monica Mountains, or how the age of buildings affects the death rates of their residents.

If Ferguson broke down his answer into many different subgroups, it would just leave us with the annoying necessity of adding those numbers up ourselves.

Rob Manzoni on January 30, 2021 7:48 AM at 7:48 am said:

Not revealing the calculations – or the code used in a prediction – is the ploy of a typical fake scientist.
Ref: Michael Mann, Phil Jones and the cabal of liars who enable idiot Nobel Laureates like Al Gore and Barak Obama to scare children; and to dismantle real science and very necessary proper education…

Nick Keefer on May 16, 2020 6:50 AM at 6:50 am said:

X's and Y's? The statistics from Sweden and other non lock down country's have proven unequivocally that his he and his model's were wrong. Once again. He has shown he could give a shit about the humans he supposedly works to protect. He sold out to big Pharma. Why so much support for this clown?

Reply ↓

Basil on May 16, 2020 7:02 PM at 7:02 pm said:

Can you read? "In the (unlikely) absence of any control measures or spontaneous changes in individual behaviour […]". In Sweden this absence is clearly not the case, as proven by their similarly troubled economy. If X is not true, not Y does not disprove X -> Y.

Reply ↓

confused on May 17, 2020 2:16 AM at 2:16 am said:

I think the original model suggested that the deaths in a scenario with some mitigation but not lockdowns would be half the totally-unmitigated scenario (1.1 million for the US, vs. 2.2 million totally-unmitigated).

I agree the totally-unmitigated scenario was never going to happen.

But 1.1 million deaths is ~0.35% of the US population. Places that didn't lock down and did more mild mitigation don't seem to be showing that.

0.35% of the Swedish population is about 36,000 deaths, which Sweden doesn't seem to be on a trajectory to get near (they are at about 1/10 of that now).

South Dakota would expect about 3,000 deaths. They have 44 now, and their cases don't seem to be exploding (they have essentially one large cluster, in Sioux Falls, associated with a meatpacking plant).

Arkansas is similar. Etc.

0.35% of the population dying in a limited, largely voluntary mitigation scenario seems pretty reasonable for dense places like New York City/New Jersey, Lombardy, Madrid, etc. given how badly they actually were affected given much stricter measures. But it doesn't seem to translate well to other areas, including much of the US.

Reply ↓

Aaron Hall on **May 8, 2020 10:00 AM at 10:00 am** said:

Exactly so. If we have extremely successful epidemiology we won't know it, and every epidemiologist will have a "Chicken Little" reputation.

Now that may come at a very high economic price, and we want to minimize that as well.

How do we know the tradeoff? "I saved 150k lives, that's 7-8 million as per the UMB, so I saved us over a trillion dollars in human capital at cost of X…"

But we won't believe it unless you show us the code and it withstands third party scrutiny.

Reply ↓

Twain on **May 8, 2020 10:25 AM at 10:25 am** said:

Aaron,

> If we have extremely successful epidemiology we won't know it, and every epidemiologist will have a "Chicken Little" reputation.

Not necessarily. What if said epidemiologists produce rigorous data, both modeled and empirical, supporting their decisions? True, we would never know for certain if their recommended actions/policies were necessary (we cannot measure the counterfactual) — but rigorous data would stop many from crying "Chicken Little".

The problem is that some epidemiologists, like Ferguson, fail to produce the information necessary to determine if their data is rigorous (code; underlying assumptions; etc.) — which inclines some people to suspicion.

Reply ↓

Joshua on **May 8, 2020 10:47 AM at 10:47 am** said:

Twain –

> What if said epidemiologists produce rigorous data, both modeled and empirical, supporting their decisions?

The problem there is that the judgement you're implying is necessarily subjective, at least to some degree – even in an expert or technically sophisticated community, let alone among the general public.

Reply ↓

Twain on **May 8, 2020 11:10 AM at 11:10 am** said:

Joshua,

"Rigor" will always be subjective; perhaps "extensive" or "thorough" are better adjectives, but I'm not sure.

To a degree, everything criteria is "subjective". But I think there is a (relatively) clear distinction between non-rigorous, haphazard data and rigorous, quality data.

Reply ↓

Carlos Ungil on **May 8, 2020 12:41 PM at 12:41 pm** said:

Have you tried to read their reports or is your suspicion based just on what you see in the press?

Reply ↓

yyw on May 8, 2020 11:39 AM at 11:39 am said:

Extremely successful epidemiology would provide projections for a number policy scenarios, so that there is a semblance of a validation feedback loop.

Reply ↓

Joseph Candelora on May 8, 2020 12:38 PM at 12:38 pm said:

Presumably 8 policy scenarios, with a subset of them each broken into 12 subpolicies based on mitigation on/off triggers, was enough for you.

https://spiral.imperial.ac.uk:8443/bitstream/10044/1/77482/14/2020-03-16-COVID19-Report-9.pdf

Or did you want more?

Reply ↓

yyw on May 8, 2020 1:08 PM at 1:08 pm said:

Thanks for the link. Why complain about the "chicken little" reputation in that case? Just point to the models with the relevant policy scenarios and compare the predictions with reality.

Reply ↓

Joseph Candelora on May 8, 2020 8:45 PM at 8:45 pm said:

Why criticize a paper you haven't read?

Reply ↓

yyw on May 8, 2020 11:39 PM at 11:39 pm said:

I was replying to Aaron's comment about "extremely successful epidemiologist". Where in my comment above did I criticize Ferguson's paper specifically?

Not that I don't have criticism now that I've read the paper.

dhogaza on May 8, 2020 10:16 PM at 10:16 pm said:

"Just point to the models with the relevant policy scenarios and compare the predictions with reality."

Do you think the people who are screaming chicken little will actually care about this? Are you aware that after publishing the paper with scenarios, and after the government of the UK decided to take firm action after all, and after Ferguson pointed out that their projection was for 20K deaths that he was accused of having modified his model? People pointed out that the various scenarios from the beginning until they were blue in the face, but the message most of the public paying attention would've seen was that Ferguson's original 500K estimate was changed by him to 20K. Not that the scenario projections had been made and presented at once, and that the model had not been changed, blah blah.

Reply ↓

Rob A on September 14, 2020 10:40 AM at 10:40 am said:

It seems that Ferguson made some pretty bad predictions about countries that didn't lock down, eg, Sweden and Japan.

https://www.iedm.org/the-flawed-covid-19-model-that-locked-down-canada/

Reply ↓

FabianB on January 17, 2021 8:33 PM at 8:33 pm said:

I would say that it looks like this free-market thin tank op-ed from June 2020 didn't age well, but it was never good.

Reply ↓

Tim on September 26, 2020 4:56 AM at 4:56 am said:

Saying "if you don't act, X will (or at least may) happen" is not a wrong prediction if people take action to avoid X, and then X fails to happen.

That's a specious argument, as it rests on the assumption that the first claim is an absolute truth, but if it isn't then the rest falls (and the use of 'may happen' is a total get out clause).

In the case of coronavirus, one only has to look at Sweden. If we follow Professor Ferguson's assertion for the UK that if we don't lockdown, up to 500,000 people could die, then a similar cataclysmic outcome should have affected Sweden, where they didn't go into lockdown. It didn't happen, and indeed thus far they have done better than we have in per capita terms. The virus did get into care homes in Sweden, as it did in many countries, but that is the only failure really.

Reply ↓

Martha (Smith) on September 26, 2020 5:54 PM at 5:54 pm said:

Tim said,
"If we follow Professor Ferguson's assertion for the UK that if we don't lockdown, up to 500,000 people could die, then a similar cataclysmic outcome should have affected Sweden, where they didn't go into lockdown."

This reasoning is not valid unless other relevant conditions were the same for Sweden then as for the UK now. Do you have any evidence that these other relevant conditions were indeed the same for Sweden then as for the UK now? (Also, note the difference between "could die" and "should have affected Sweden" — the first describes a possibility, the second a certainty.)

Reply ↓

Tracey Hoolachan on December 13, 2020 1:40 PM at 1:40 pm said:

This is a typical con trick.
It was never a bad model from the start only something that needs adjusting. Health will always need adjusting because our lives are guided by trends, govt budgets & the labour market that are always changing. Those trends in food driven by marketing increased sugar increasing a chronic disease diabetes, increased computer use over sport lower vitamin D absorption. That govt austerity malnutrition & homelessness, reduced free education so less medics. Labor markets driven by industrial/robotic shifts unemployment leading to poverty.
Even brainwashed by Behavioural Insight Team the Covid dreams & hair loss tells me people will always be individuals, will not blindly accept being told what they have to do. They banded together against them and more are waking up every day and joining.

Reply ↓

Tom on May 8, 2020 9:40 AM at 9:40 am said:

Mad Cow paper:
https://www.ncbi.nlm.nih.gov/pubmed/11786878

From abstract:

"Extending the analysis to consider absolute risk, we estimate the 95% confidence interval for future vCJD mortality to be 50 to 50,000 human deaths considering exposure to bovine BSE alone, with the upper bound increasing to 150,000 once we include exposure from the worst-case ovine BSE scenario examined."

Consistent with the "up to 50,000" quote but the quote fails to mention the lower bound.

Reply ↓

Dale Lehman on May 8, 2020 9:46 AM at 9:46 am said:

Thank you for saving me a bunch of time trying to track down this information. I will now paraphrase Andrew's question: how should we judge the media? I am willing to ignore this journalist's work entirely now – as evidence by his worst work. He has clearly misstated Ferguson's BSE work and that is sufficient for me to ignore the rest of his (Fund's) work. Everyone can make mistakes, including journalists, but this one looks dishonest and my patience of dishonesty is rapidly disappearing.

Reply ↓

dhogaza on May 8, 2020 12:14 PM at 12:14 pm said:

"I am willing to ignore this journalist's work entirely now – as evidence by his worst work."

If he's the John Fund I found through google, though described as being a journalist, he's paid to write editorial pieces for two conservative American outlets, The National Review and The American Spectator. No journalistic rigor is asked for or expected. He writes to support the political ideology those outlets were created to support, hell or high water.

Reply ↓

ira on May 8, 2020 4:20 PM at 4:20 pm said:

Exactly. If you know who John Fund is there's really no reason to read any more.

Reply ↓

Kaya3 on May 10, 2020 6:26 PM at 6:26 pm said:

Hilarious watching people dismiss something because it doesn't comply with their views. The whole of mainstream media is corrupted to it's core. If the content is wrong then show it.
Ferguson is a clown. An idiot who has been wrong almost every time he opened his mouth and has caused immeasurable damage to the whole planet and the deaths of tens of thousands who will die as a result of the actions taken over this event.
Yes he left out the lower parameter, so what? Any scientist who gives themselves a range of 50 to 150,000 should never be listened to in the first place. He's a moron, and now he's resigned because he couldn't even follow the orders he was imposing on the rest of the population to shag his married lover. Go figure.

Reply ↓

Carl on March 25, 2021 6:40 PM at 6:40 pm said:

You can say the same thing of the "journalists" who write for liberal and leftist outlets.

Reply ↓

Joshua on May 8, 2020 10:48 AM at 10:48 am said:

Tom –

> Consistent with the "up to 50,000" quote but the quote fails to mention the lower bound

Bingo.

Reply ↓

Andrew on May 8, 2020 11:03 AM at 11:03 am said:

Tom:

Thanks for looking this up. That's helpful. I should've thought of googling *Neil Ferguson BSE sheep* or something like that before pressing Send on my post!

Reply ↓

Adede on May 8, 2020 11:49 AM at 11:49 am said:

50 to 50,000 is such a huge range. Almost anything would fall into it.

Reply ↓

jd on May 8, 2020 11:57 AM at 11:57 am said:

Came across this blog post that I thought was pretty good: https://maybury.ca/the-reformed-physicist/2020/04/09/no-better-than-a-fermi-estimate/
Especially the last paragraph, which I quote here:
"During World War II, [Nobel laureate, Ken] Arrow was assigned to a team of statisticians to produce long-range weather forecasts. After a time, Arrow and his team determined that their forecasts were not much better than pulling predictions out of a hat. They wrote their superiors, asking to be relieved of the duty. They received the following reply, and I quote, "The Commanding General is well aware that the forecasts are no good. However, he needs them for planning purposes."

Reply ↓

Martha (Smith) on May 8, 2020 3:44 PM at 3:44 pm said:

!

Reply ↓

TJ Raitzer on May 8, 2020 10:31 PM at 10:31 pm said:

They should have then asked for a better hat…

The entire post is very much worth reading, especially "there's no need to solve differential equations" when "you can get these results by counting on your fingers…"

Reply ↓

Phil on May 8, 2020 12:39 PM at 12:39 pm said:

Presumably the 'central' estimate was 5000, with factor of 100 uncertainty either way. I agree it's a huge range.

But (1) if you really are that uncertain, you should say so. That's hard for people to do, so I see this as a feature, not a bug, unless there was some reason to think he should have narrowed the bounds.

(2) it's tempting to say there's nearly no information in an estimate with bounds this wide, but there is. Contrast it with an estimate of 50-5000. If you give an estimate like that you're saying even the worst case isn't that bad. Put the upper limit at 50K and you're saying Jeez, we have no idea. The degree of ignorance is the point, if you give bounds that wide, and that's a kind of information.

Reply ↓

Bill Spight on May 8, 2020 5:18 PM at 5:18 pm said:

Yes, you should say how uncertain you are, if you can. That may be hard for people in general to admit, but it's a sine qua non for a scientist to say, especially one with even a smidgeon of statistical training.

Consider the Arizona epidemiologist who simply said that unless Arizonans continued to shelter until the end of May, he did not see how to avoid exponential growth in Arizona cases. His statement was relevant, important, and appropriately vague.

Reply ↓

Zhou Fang on May 8, 2020 2:02 PM at 2:02 pm said:

The death toll of Covid 19 does not...

Reply ↓

More Anonymous on May 8, 2020 2:51 PM at 2:51 pm said:

Adede: Wide ranges show up like this because, at the start, there often isn't enough evidence to tell if an epidemic will take off and cause many deaths or die out and cause a small number.

Until more precise estimates of quantities like R0 can be obtained, the only honest answer is a wide range of deaths. People want smaller ranges and refusing to give them is a principled position.

When people then cherry pick the side of the range that didn't happen, I think they are being misleading and dishonest (not you, John Fund).

Reply ↓

Linch Zhang on May 9, 2020 3:05 AM at 3:05 am said:

I don't think it's that large of a range tbh. 3 orders of magnitude honestly is pretty impressive when you're considering that Ferguson

1) Was trying to predict different numbers based on whether an outbreak will happen or not, at the beginning of an outbreak.
2) Was giving 95% CIs, so if he's well-calibrated, only 5% of the time will numbers fall outside of his range.
3) Was trying to predict numbers of people who will die from a novel infectious disease *before 2080*.

If you think you can do better, you try to make a prediction about deaths from an arbitrary novel infectious disease before 2080.

Heck, I'll throw you a lowball and give you a disease that we already know a lot about. How about Ebloa?

Global deaths from Ebola before 2080. What's your 95% CI?

Reply ↓

Robert Rees on April 13, 2021 5:38 AM at 5:38 am said:

You raise an interesting point which is also pertinent to the current debate. I have noticed that the projections from these model have often extremely wide error bars. I agree that to write 'Up to 50,000' is misleading. But to project a range of 50 -150000, with a central estimate of 50,000 is effectively unfalsifiable and hence unscientific. The whole point of predictions is to prove or disprove. These type of wide ranges do neither.
It also suggest to me that the model is pretty chaotic is small changes in assumptions can lead to massive changes in outputs – as has been seen in the Covid models.

Reply ↓

BenK on May 8, 2020 9:49 AM at 9:49 am said:

You include a good panel of options. Just in January, I was with a group of people assembled to forecast biological disasters and around the coffee bar they were panning Tony Fauci for having had career success on the back of influenza predictions that were worse than the outcomes of the pandemics. These putative experts were saying that COVID-19 would be another SARS, limited to East Asia, no big deal.
They lumped him in with people like Ferguson and Marc Lipsitch, who were apparently always wrong. According to them.

Reply ↓

More Anonymous on May 8, 2020 2:06 PM at 2:06 pm said:

+1

Reply ↓

Ben Bolker on May 8, 2020 10:04 AM at 10:04 am said:

This article seems like a hatchet job. While anyone can turn out to be a charlatan, Ferguson has a very good reputation in epidemiological circles as a thoughtful, sensible guy. In support Fund quotes:

* Johann Giesecke, a respected epidemiologist but one who disagrees strongly with Ferguson (and the majority of epidemiologists) on the wisdom of lockdown-based strategies
* Elon Musk (!!)
* Jay Schnitzer, "an expert in vascular biology and a former scientific direct of the Sidney Kimmel Cancer Center in San Diego"
* Charlotte Reid, "a farmer's neighbor" (on the foot and mouth disease virus)

Everyone (including Ferguson) agrees that Ferguson behaved foolishly. It's not crazy (although I think it's incorrect) to say that Ferguson's model was wrong, and that lockdown strategies are misguided. But calling him a charlatan is just unfair.

Also, I think your title is a little clickbaity.

Reply ↓

Andrew on May 8, 2020 10:29 AM at 10:29 am said:

Ben:

Yeah, I noticed the flakiness of the sources too. Regarding clickbait, my tile is taken from the original article and I kept it in quotes, but maybe that is not clear to the reader.

Reply ↓

Ben Bolker on May 8, 2020 10:29 AM at 10:29 am said:

FWIW here's the source of the "150 million dead from avian influenza" quote
https://www.theguardian.com/world/2005/sep/30/birdflu.jamessturcke

Someone else (David Nabarro, WHO) gave a range of 5-150 million dead in an avian influenza outbreak. Ferguson said: "Around 40 million people died in 1918 Spanish flu outbreak … There are six times more people on the planet now so you could scale it up to around 200 million people probably."

A worst-case back-of-the-envelope calculation, but not at the level of a misguided prediction (IMO).

Reply ↓

yyw on May 8, 2020 11:47 AM at 11:47 am said:

From the Guardian: Last month Neil Ferguson, a professor of mathematical biology at Imperial College London, told Guardian Unlimited that up to 200 million people could be killed. "Around 40 million people died in 1918 Spanish flu outbreak," said Prof Ferguson. "There are six times more people on the planet now so you could scale it up to around 200 million people probably."

If the reporter had not misquoted/misinterpreted him, he deserves a lot of flack for this kind of stupid speculation.

Reply ↓

Anonymous on May 8, 2020 11:42 AM at 11:42 am said:

I don't know anything about Ferguson, but that Elon Musk attack "Elon Musk calls Ferguson an 'utter tool" who does "absurdly fake science."' makes me want to root for Neil.

Some amusing Musk quotes in the same vein:

http://www.elonmusk.today

Reply ↓

Andrew on May 8, 2020 11:48 AM at 11:48 am said:

Anon:

These quotes are hilarious. We can await the next article by John Fund: a debunking of Elon Musk, ending with the statement, "So the real scandal is: Why did anyone ever listen to this guy?"

On the plus side, Musk didn't call Ferguson a pedo guy.

Reply ↓

dhogaza on May 8, 2020 12:18 PM at 12:18 pm said:

Also on the plus side, if you buy one of his cars, you don't get Elon.

Reply ↓

More Anonymous on May 8, 2020 3:34 PM at 3:34 pm said:

> *We can await the next article by John Fund: a debunking of Elon Musk, ending with the statement, "So the real scandal is: Why did anyone ever listen to this guy?"*

Personally, I think John Fund should write a new article with that same conclusion every week. The last one in the series can be about himself.

Reply ↓

Martha (Smith) on May 8, 2020 3:51 PM at 3:51 pm said:

I don't know whether to say, "Ooooh!" or "Meow!" to this.

Reply ↓

Bill Spight on May 8, 2020 5:26 PM at 5:26 pm said:

Foot in mouth disease is unfortunately endemic in human populations. Even worse, it often goes undiagnosed.

Reply ↓

Rahul on May 8, 2020 10:05 AM at 10:05 am said:

Did you guys see the scathing code review of his model? My question is, where does this stand compared to the median code for models in scientific literature. From what little I have seen, a lot of scientific conclusions hinge on programs that are atrociously bad and buggy.

https://lockdownsceptics.org/code-review-of-fergusons-model/

Reply ↓

D Kane on May 8, 2020 10:17 AM at 10:17 am said:

I was about to post that. It is amazing and deserves a thread of its own. **Ferguson's sin is not accuracy of his forecasts but the quality of his code.**

> a lot of scientific conclusions hinge on programs that are atrociously bad and buggy.

"A lot", in my experience looking at academic code, is about 95%+. To anyone who comes from an industry background, it is shocking. But perhaps it shouldn't be. Code poorly in industry and your business fails. Code poorly in academia and . . . nothing much?

Reply ↓

Rahul on May 8, 2020 11:12 AM at 11:12 am said:

Exactly. I think code is bad because the incentive structure is just not there to write good code.

A android app or industrial PLC or banking website fails due to bad code (and I'm not even saying code quality is awesome there!) you have a bunch of angry customers and some monetary whipping. Real chance of losing a job.

An academic epidimiological, election or global warming model throws out crap and who notices and even if they do, what are the real consequences to anyone?

Reply ↓

jim on May 8, 2020 11:24 AM at 11:24 am said:

"the incentive structure is just not there to write good code. "

I agree but it's more than just a customer issue.

Isn't most academic research code written by a single person? For it to be so bad it almost has to be. It's basically the researcher's personal project. Under those circumstances, it's really not worth the time to properly organize and structure it, since the person using it mostly knows what's going on.

The code doesn't belong to an organization. It belongs to an individual. If that person changes jobs s/he takes the code with. So it's not necessary for the code to be accessible to others for the purpose of perpetuating the project.

Reply ↓

Rahul on May 8, 2020 12:45 PM at 12:45 pm said:

@jim

I agree somewhat about maintainability. But the real issue is quality control. In a one man project you do you ensure that the results are *right*?

Almost axiomatically, if the result is something critical or that we care about then relying on the results of a poorly written, unaudited, one man code would be foolish. It continues because the results dont mean much beyond publication and presentations.

At the very least, I would independently fund multiple one-man-codes and cross check results.

Reply ↓

Daniel Lakeland on May 8, 2020 1:10 PM at 1:10 pm said:

A virus is doubling every 3 days in a population, you have about 9 days to make a decision before the naive extrapolation says it's too late and your hospitals will be overwhelmed. Do you

1) Spend 60 days code reviewing some code, and running an extensive set of test cases, then write a paper, wait for it to be peer reviewed by a group of people who will take 3 months, and publish it at that point?

2) Rely on whatever quality control you did years ago, which was probably fairly extensive but lets be honest no one is really *sure* that the models are coded correctly, but they do qualitative produce the results that are expected and agree qualitatively with other models?

3) Let Boris Johnson make the decision based on his "gut instinct" alone even though he has *exactly zero* knowledge of infectious diseases?

It's obvious (2) is the only answer available. That it's not perfect isn't the issue. it's the best you've got.

Reply ↓

Rahul on May 8, 2020 1:17 PM at 1:17 pm said:

@Daniel:

I am criticizing the general practice. In academia in "peacetime" too with no virus chasing you the code quality is atrocious is the premise I have. And I was one of those coders too during my PhD.

In fact, I am defending Ferguson in some sense: I say he's no worse than the typical academic code that gets you published.

I'm saying we just have a culture of crappy codes in academia. Simply because the outcomes predicted are mostly so low-stakes that people just don't care!

SDE on May 8, 2020 1:22 PM at 1:22 pm said:

I don't know . . . It seems there is a fourth option:
4) Recognize that person to person transmission is happening for a disease that obviously is serious, and social distancing has been used effectively since before the germ theory of disease, acknowledge the math of exponential growth, and use that to drive a decision. And all the time, don't rely on an implementation of a model that is fatally bad.

In this case, basic knowledge of epidemic mechanism, and some charts can drive the decision without the noise of a model implementation that is deeply, deeply flawed and in unpredictable ways. Again, this is based on my own reading of the code, not some blog post.

Daniel Lakeland on May 8, 2020 2:22 PM at 2:22 pm said:

SDE: I agree with you. My own take on this model output was that it was just "drawing some charts" to show *qualitative* differences in outcomes. There are massive, order-of-magnitude changes caused by quarantines, therefore the right solution was quarantine until more information is available to make more informed decisions.

The model did its job of communicating in a viscerally understandable way the dramatic differences that quarantine can bring about.

If someone had just “drawn some infovis” on a flip chart, I doubt we would have seen implementation of effective quarantines.

From a political perspective, the model did its job of convincing people to take *very basic ideas* seriously because without complex blinkenleitz they were going to blow through the population for herd immunity and get hundreds of thousands of deaths or more, and react too late to do anything about it.

Phil on **May 8, 2020 1:09 PM at 1:09 pm** said:

When I was in a quasi-academic setting (Lawrence Berkeley National Laboratory) I wrote pretty bad code. Writing good code is hard and takes more time, or at least it does for me. The premium for writing good vs bad code has been decreasing as I’ve been putting more effort into writing good code, but it’s still substantial.

As a researcher, most of these were true for any single project.
– My code only had to work with the data I had in hand. I could put in kludgy code to handle missing data or corner cases or whatever, without worrying about what would happen with other data.
– No one except me would ever see the code. I only had to leave myself enough documentation to remember why I had coded something a particular way, not to make it clear to other people.
– Output could usually be an R object, which I would use directly for preparing figures and tables for publication. No fancy (or non fancy) interface for I/o.
– Checking to make sure the code did what it was supposed to could be ad hoc. Feed fake data with a known ‘signal’ into a function , check the output and make sure it looked about like I expected, do it a couple more times, great, call it good. I might not even note that I had done this check.

…all of which was somewhat embarrassing when people asked for my code, which did sometimes happen (and I would send it). But I found that people did not mock me for my poor code, and it really was poor. It turns out people are used to seeing bad code, there’s lots of it out there.

Now, I work as part of a small team— two or three of us — and although we divvy up the analyses that we do, we do have to use or alter each other’s code sometimes. I’ve learned a lot from seeing the code my partners write and it has made me a much better programmer. My code is easier to understand , easier to check, just better all ways than what I used to write.

Reply ↓

Daniel Lakeland on **May 8, 2020 1:21 PM at 1:21 pm** said:

Phil, in addition to all this I’ll add that making your code modular and capable of being used in a wide variety of ways is frequently a *terrible* idea. You have to think about the probability that it *will* be used in such a wide variety of ways, and further that the choices you made about how to modularize it will turn out to be the correct choices under the real-world scenario when you actually do try to utilize it.

Some things are relatively easy to see how to modularize, but often it’s impossible to foresee the use your code might be put to in the future.

It’s much more important to handle all the corner cases in your *actual* problem so you get the right answer in that problem than the make it easy and quick and low cost to get the *wrong* answer to a lot of theoretical future problems.

This is particularly the case when it comes to dealing with external inputs. If someone hands you a dataset where missing values are sometimes NA, sometimes N/A, sometimes NULL, sometimes 0, sometimes -1, sometimes -999, sometimes "0", sometimes "not avail" sometimes "Not available"… you had better take care of all those cases, and there's no reason to think that what you do here is generalizable in any way.

All that being said, the more you know and practice general code quality issues, the easier it becomes to do a good job the first time, and get the realistically good balance between understandable code and handling all the corner cases.

Reply ↓

Rahul on May 8, 2020 1:27 PM at 1:27 pm said:

Some of these changes need to be systematic and starting with funding culture:

e.g. Funding agency sets a target that if I spend say $100,000 funding a PhD student writing a code I will keep aside $25,000 to fund another (independent) student to vet the code or at least run test cases etc. to validate the code.

Anecdote: I had to use a Density Functional Theory (QM) code to predict material properties as a student. I proposed, shall we spend the first few months validating predictions of this code for known materials? Professor leading project says "Waste of time"

Reply ↓

Daniel Lakeland on May 8, 2020 2:33 PM at 2:33 pm said:

Agreed. That the Stan group for example has consistently prioritized code quality, testing, and validation has been a HUGE win for the world. The same is true for other groups where that's been the case. LAMMPS for example is an excellent molecular dynamics / mechanics code. The effort has been put in to maintain it, squash bugs, and etc.

The Julia language project is building a modern language for scientific computing that is basically a LISP with a scientific computing syntax and a strong emphasis on type-inference leading to high speeds. It displays standards that are way way better than typical industry standards.

That there isn't a high quality off-the-shelf infectious disease *modular modeling platform* built most likely by a partnership between the CDC, the NIH, the NSF, the European CDC / WHO, and tens to hundreds of universities over the last 20 years is *atrocious* but it's not the fault of people like Ferguson. It's the fault of priorities at the CDC, NIH, NSF, WHO etc for failing to prioritize such a thing.

dhogaza on May 8, 2020 12:21 PM at 12:21 pm said:

"Code poorly in industry and your business fails"

You mean just like when MicroSoft went bankrupt as a result of years of refusing to address a seemingly endless series of brain-dead bugs in their software?

Reply ↓

Rahul on May 8, 2020 12:48 PM at 12:48 pm said:

It's all relative: Sure, compared to the Linux kernel MS-code could be a mess. But compared to most one-man academic software MS has built the equivalent of the Hoover Dam.

Don't get me wrong: There's great academic coding projects but those are mostly well maintained, multi year open source teams. Stan may be a good example.

Reply ↓

Daniel Lakeland on May 8, 2020 1:28 PM at 1:28 pm said:

> compared to most one-man academic software

so you're saying that an enterprise employing around 150,000 people spending a large fraction of its $140B/yr of revenue on continuously attempting to improve its shitty code over a period of 40 years is doing substantially better but still not great compared to one or a handful of guys spending a few thousand hours over a few years all around a decade or more ago with a total budget of probably substantially less than $200k

Reply ↓

Rahul on May 8, 2020 1:48 PM at 1:48 pm said:

No. I am saying we should be improving things at the bottom especially when we fund academic software development. Put a premium on independent code validation and transparency and better coding standards.

Microsoft has the Market to whip it, at least to the extent markets function. There's a reason why Linux took over so much of the enterprise server market. Or similarly for the successes of Chrome, or Apache etc.

Problem with academic code of this sort is that there is not much of a market. Most codes are written, tested and results published by the same one Professor. Quality of prediction or code rarely has much of an impact.

Reply ↓

Daniel Lakeland on May 8, 2020 2:08 PM at 2:08 pm said:

Indeed, I agree with the idea of "pay people to write good code, and then externally validate it". We should do a lot more of that, and also a lot more of "pay people to create curated public datasets addressing important questions" and a lot less of "pay people to pull the slot machine lever and assert causality after p is less than 0.05"

dhogaza on May 8, 2020 2:42 PM at 2:42 pm said:

"There's a reason why Linux took over so much of the enterprise server market. Or similarly for the successes of Chrome, or Apache etc."

None of which come close to meeting the standards the self-appointed gods of code review that have mysteriously popped up here recently claim are necessary for the production of reliable code.

I don't argue that academic code shouldn't be better (or industry code, for that matter, most of which is arguably shit, except for mine, of course :))

What I would argue is that the folks who've showed up here pontificating about how horrible Ferguson's code is have no real interest in whether or not it produces reasonable results.

Their motivation is different.

Rahul on May 8, 2020 3:14 PM at 3:14 pm said:

A good story about a bug in academic simulation code:

The war over supercooled water
How a hidden coding error fueled a seven-year dispute between two of condensed matter's top theorists.

https://physicstoday.scitation.org/do/10.1063/PT.6.1.20180822a/full/

To me this is just the tip of the iceberg. If anyone started validating published articles in the simulations & models area I think the dominoes will topple all over the place. We really need an Ioannidis here!

Reply ↓

Ben on May 8, 2020 1:21 PM at 1:21 pm said:

> Ferguson's sin is not accuracy of his forecasts but the quality of his code.

Sure software engineering is a factor, but it's hardly the only thing.

I've heard a lot about Ferguson's code but we've yet to discuss his model. Presumably if it's a bad model, it doesn't matter how much software engineering you throw at it. Judging by the conversation, the issue is the model isn't that clearly defined, and that is reflected in the code.

Ofc. I know what I know about Ferguson from reading this blog, so I could be wrong and his model could be excellent, just poorly coded.

> Code poorly in industry and your business fails

https://www.safetyresearch.net/blog/articles/toyota-unintended-acceleration-and-big-bowl-%E2%80%9Cspaghetti%E2%80%9D-code

Toyota is still rockin' along as far as I know.

Reply ↓

dhogaza on May 8, 2020 2:44 PM at 2:44 pm said:

Ben:

"I've heard a lot about Ferguson's code but we've yet to discuss his model. Presumably if it's a bad model, it doesn't matter how much software engineering you throw at it."

Bingo! I don't know if you wrote the above before or after I wrote something similar elsewhere in this steadily lengthening set of responses to the OP, but yeah. Odd, isn't it? (actually, no, it isn't odd, deflecting from the actual model is the point.)

Reply ↓

Ben on May 8, 2020 6:25 PM at 6:25 pm said:

I did see it! And thanks for all your posting today (https://statmodeling.stat.columbia.edu/2020/05/08/so-the-real-scandal-is-why-did-anyone-ever-listen-to-this-guy/#comment-1332267).

Reply ↓

J on May 8, 2020 8:52 PM at 8:52 pm said:

Not true. Awful code is common in industry also.

Reply ↓

Simon on May 8, 2020 10:18 AM at 10:18 am said:

The url hints at more than a little bias from the "reviewers" plus the code they are looking at isn't that used for the simulations rather the "cleaned up" code – a process that is likely to have introduced bugs.

Reply ↓

Rahul on May 8, 2020 11:14 AM at 11:14 am said:

Sort of stupid of Ferguson to release "cleaned up" code that isn't the one generating the actual conclusions bandied about.

At the very least, wasteful of debuggers time. Defeats the whole purpose of code review and posting code online.

Reply ↓

jim on May 8, 2020 11:26 AM at 11:26 am said:

"Defeats the whole purpose of code review"

+1

Reply ↓

John Fund on May 8, 2020 12:48 PM at 12:48 pm said:

I am John Fund, the author of the article being discussed here, I couldn't fully explain every objection to Ferguson due to an editor's word limit in that format. But pay close attention to what I did include in the excerpt above: "Last March, Ferguson admitted that his Imperial College model of the COVID-19 disease was based on undocumented, 13-year-old computer code that was intended to be used for a feared influenza pandemic, rather than a coronavirus. Ferguson declined to release his original code so other scientists could check his results. He only released a heavily revised set of code last week, after a six-week delay."

This lack of transparency has NOT been refuted, the original code isn't available and won't be. This is why several people in the field were willing to talk to me about Ferguson and why they harbor suspicions about transparency in this and his other models.

Reply ↓

Andrew on **May 8, 2020 1:01 PM at 1:01 pm** said:

John:

Thanks for responding. It seems like the real concern here is the specifics of the model and the code; the criticism of Ferguson's past projections regarding mad cow disease etc., is a red herring, as it's not clear that those projections were in error. Perhaps the right question to ask is not, "Why did anyone ever listen to Ferguson," but, rather, "Were Ferguson's projections given undue influence in government decisions?"

Reply ↓

Dale Lehman on **May 8, 2020 1:08 PM at 1:08 pm** said:

Not for me, Andrew – and John. You have severely mis-characterized Ferguson's past predictions – and now you have chosen not to respond to 2 instances where this has been pointed out. I frankly don't care at this point about whether Ferguson's credibility should be questioned. It is your credibility that concerns me.

Reply ↓

John Fund on **May 8, 2020 1:12 PM at 1:12 pm** said:

Fair point. In mid-March, Britain's government was debating strategy on the virus. Ferguson had notable seat at table. His model had range of projections, but from previous experience in government he should have known the politicians would pay most attention to the most perilous worst-case scenarios. From what I've been told, he encouraged focus on those scenarios (called tilting the argument) as he had in previous predictions.
I blame the politicians most – someone should have said: "How close have the upper-limit projections been to being accurate in previous crises? Shouldn't we factor that into our decision?" Instead, the politicians largely panicked (several including the PM ironically came down with COVID soon after).

They went for what everyone acknowledges was the first time in history a healthy population in Britain was imprisoned instead of a. focus on protecting the old, vulnerable and frail (see lack of attention to nursing home issues). Can't find anyone in the mid-March debate who said that lockdowns had scientific evidence behind them – other than the Chinese unreliable numbers.

I think politicians panicked in Britain – and some scientists enabled them – as they had in the past.

Reply ↓

dhogaza on **May 8, 2020 1:22 PM at 1:22 pm** said:

Right-wing political hack John Fund says:

"other than the Chinese unreliable numbers"

Ummm actually the Chinese numbers have held up quite well. Of course, they really DID imprison the residents of Wuhan, thus their success.

"Can't find anyone in the mid-March debate who said that lockdowns had scientific evidence behind them"

Someone needs to help you improve your google-fu. Perhaps you need to search backwards in time before this current epidemic. It's not like this is new stuff.

And the UK has not "imprisoned" its populace. The kind of rhetoric undoubtably plays well to your usual audience, but will quite likely lead to most people here not taking you very seriously.

But to borrow your terminology:

Ferguson's model projected 7,000-20,000 deaths in the UK "imprisonment" scenario.

There have been 35K UK deaths thus far.

So yeah, it's done a poor job. But its lousy projection for the scenario that was implemented doesn't exactly support your belief that the lockdown was unnecessary.

You forgot to mention "but, Sweden!!!!"

Anonymous on May 8, 2020 2:03 PM at 2:03 pm said:

"Ummm actually the Chinese numbers have held up quite well."

Just for the record dhogaza, you think China's claim of fewer than 5000 deaths is accurate?

Ben on May 8, 2020 2:28 PM at 2:28 pm said:

> several including the PM ironically came down with COVID soon after

I don't understand how that is ironic.

Zhou Fang on May 8, 2020 2:32 PM at 2:32 pm said:

5000 deaths is an underestimate yes, but the progression of the Wuhan pandemic is likely qualitatively accurate in the sense of the shape of the curve and the end of the epidemic following lockdown.

If we merely import over the figures from other countries, we would expect that case numbers and death numbers are probably undercount by perhaps orders of magnitude. However, the failure of Beijing and other cities to look like Wuhan, the CCP's willingness to end Wuhan's lockdown, the lack of exported cases from other Chinese cities, the genotype record, all of that suggests that the overall picture of a successful lockdown seems accurate.

dhogaza on May 8, 2020 2:47 PM at 2:47 pm said:

"Just for the record dhogaza, you think China's claim of fewer than 5000 deaths is accurate?"

What Zhou Fang said, but better than I would've.

Anonymous on May 8, 2020 3:06 PM at 3:06 pm said:

dhogaza,

So “held up pretty well” means “off by orders of magnitude”.

For the record I think that’s right. The only thing remaining to determine is whether the order of magnitude they’re off by is 2 or 3.

dhogaza on May 8, 2020 3:15 PM at 3:15 pm said:

Ugh, Zhou and I have fallen for a deflection:

“other than the Chinese unreliable numbers”

Let’s run some more recent numbers.

Antibody testing results in NY lead to an infected fatality rate of about 0.8%, which fits reasonably well within the original WHO estimate of 0.3%-1% they published after sending a team to China way back in February. There are all sorts of issues with these serology tests, but at least the NY test had enough positive results to minimize issues due to false positives.

R0 seems to be thought to be around 3. leading to about 66% of the population being infected before herd immunity kicks in.

So absent mitigation efforts (voluntary or government imposed) we get

326,000,000 x 0.66 x 0.008 as a reasonable estimate of the number that might die in the US if we let it run to herd immunity, which was essentially Boris Johnson proposed as policy in the UK in the beginning.

The above gives about 1.7M deaths using the NYC serology study results. Ferguson’s model “up to 2.2 million”, with a lot less available data to work with.

We could plug in UK population data, but the comparison won’t differ.

“I think politicians panicked in Britain – and some scientists enabled them – as they had in the past.”

IIRC Ferguson said “up to 500K” might die in the UK if no action were taken. 77% of that is 386K.

Would the UK government have acted differently if the number were 386K rather than 500K? Remember, 386K doesn’t depend at all on Chinese data.

Zhou Fang on May 8, 2020 3:21 PM at 3:21 pm said:

They can be off by a proportional quantity and still validate the core issue at hand. Indeed if there were hundreds of thousands dead in Wuhan instead of thousands, then the fact that we didn’t see comparable numbers dying in Beijing and Shanghai and every other city right now (which we do not need China’s figures to conclude) would strengthen the argument for lockdown, not weaken it. If you want to claim that China’s real figures would invalidate lockdowns, then you would want to claim that China *overstated* their covid numbers, not understated it.

Joseph Candelora on May 9, 2020 2:33 AM at

2:33 am said:

John Fund,

You say "From what I've been told, he encouraged focus on those scenarios (called tilting the argument) as he had in previous predictions."

It is wholly inappropriate for a journalist to make such unsourced claim. What exactly were you told? By whom?

And you snuck in "as he had in previous predictions", which can only be reference to the claims elsewhere
in your article that he has provided bad projections in other epidemics. Claims of yours whose fraudulent nature is easily uncovered with even a brief glance at the primary sources, and which have been well debunked by other commenters upthread*.

Additionally, I question your understanding of the March 16 ICL report. Exactly what "most perilous worst-case scenarios" are you referring to? If you mean his projection of the death toll if zero mitigation were undertaken (voluntarily or by mandate), then of course that got significant attention. It demanded attention. But it was simply the baseline, his model's best estimate of what would happen if nothing was done and a basis against which to measure the estimated effects of mitigation. And frankly, there's very little that's controversial about those projections. His model resulted in a 0.9% infection fatality rate, and that with zero mitigation the virus would not be eradicated until it had spread to 80% of the UK population. Neither of these can be known, of course, but NYC fatality data supports the 0.9% IFR as reasonable, and 80% infected is a reasonable outcome based on the transmissibility of the virus. There is nothing then extreme or "worst case" about his model's projection of 510k deaths unmitigated; just a dumb application of those numbers to UK's 66m population gives 66m * 80% * 0.9% = 475k deaths — pretty well in line.

If you want to talk worst-case scenario, that would've been using the high end of the 95% confidence interval for IFR they saw from their model — which was 1.4% and would've increased unmitigated deaths to ~800k in the UK. That was a number they could have published in this paper, but chose not to. So please stop peddling your false narrative about upper-limit projections.

*But just to debunk a little more, read the press release for Ferguson's 2002 paper on possible human health risk of BSE infection in sheep. The paper was clearly and explicitly a call for more data, saying that the range of possible cow-based BSE deaths was from 50 to 50,000. And then For you to give only the top ends of those ranges, and to further suggest that those ranges were meant to drive decision-making rather than to simply serve to clarify data needs is the worst kind of misrepresentation.

Stephen M St Onge on May 11, 2020 7:40 PM at 7:40 pm said:

An engineer made a good point on Twitter the other day. He said, approximately, "If the model shows a low-probability catastrophic outcome, I don't say 'Maybe it' over-pessimistic, I'll ignore that possibility.' What I do is build in safety factors that can handle that outcome."

The Tacoma Narrows Bridge Collapse, the Pearl Harbor attack, the Comet passenger jet fiasco, the current Boeing problems with the 737 Max, they all seem to have been caused by deciding to exclude the unlikely possibilities as too improbable and expensive to bother with.

Basically, you seem to be saying the politicians should have ignored the possibility of disaster in setting policy. That strikes me as extremely unwise.

Bob on May 8, 2020 1:21 PM at 1:21 pm said:

Yes, given the discussions here in the past on replication, is the Imperial Covid model important enough to qualify for these transparency requirements?

Reply ↓

John Fund on May 8, 2020 1:41 PM at 1:41 pm said:

To Dhogaza: Your first point is a completely ad hominem personal smear on me. Not impressive for this high-tone forum at all. Also a form of insult I don't see anyone else on this forum using. Why are you the outlier in behavior on this forum?

"Chinese numbers have held up quite well." Hmm…really? .I suggest you might do a bit more research. One of MANY troubling examples. if the numbers really hold up, I await the Chinese allowing in international experts to investigate and speak freely to medical personnel:
" In northwest China's region of Xinjiang, authorities claim that by April 5, a total of only 76 people have contracted the disease, and only three have died, in a population of roughly 24.5 million. On the same day, the United States reported a total of 304,826 cases and 7,616 deaths among 331 million people. In other words, Beijing claims that, on a per capita basis, roughly 300 times as many people have contracted the disease, and more than 180 times as many have died, in the United States than in Xinjiang. Everything is possible. But claiming that the region where an estimated 1 million Muslims are imprisoned in concentration camps is far safer from the coronavirus than the United States strains credulity."

Bob on May 8, 2020 1:53 PM at 1:53 pm said:

"The Chinese numbers have help up quite well".

You've said some foolish things on this thread, that takes the biscuit.

Zhou Fang on May 8, 2020 2:36 PM at 2:36 pm said:

John:

Why? The infection is a function of access. If the infection doesn't get into the camps – and those camps are kept remote and isolated, why does it matter how many people are in them? The failure to close down transport from New York, for instance, makes the US far more vulnerable to covid spread than China.

spread than China.

dhogaza on May 8, 2020 2:51 PM at 2:51 pm said:

"But claiming that the region where an estimated 1 million Muslims are imprisoned in concentration camps is far safer from the coronavirus than the United States strains credulity."

Again, what Zhou Fang said.

"strains credulity."

Not only are you arguing a logical fallacy (argument from incredulity), but you even used the "credulity" word. Thanks for not trying to hide it.

somebody on May 8, 2020 3:27 PM at 3:27 pm said:

To John Fund

I agree that starting a statement by saying "Right wing political hack John fund" is pointlessly rude. That said...

> " In northwest China's region of Xinjiang, authorities claim that by April 5, a total of only 76 people have contracted the disease, and only three have died, in a population of roughly 24.5 million. On the same day, the United States reported a total of 304,826 cases and 7,616 deaths among 331 million people. In other words, Beijing claims that, on a per capita basis, roughly 300 times as many people have contracted the disease, and more than 180 times as many have died, in the United States than in Xinjiang. Everything is possible. But claiming that the region where an estimated 1 million Muslims are imprisoned in concentration camps is far safer from the coronavirus than the United States strains credulity."

This is a truly silly take. First of all, the U.S. has over 300 times as many people as Xinjiang. More than 180 times as many dying in the United States during its infection peak than in a population 12 times smaller during its infection tail seems completely plausible. You're seeming seems to operate on a logic of "worse places will be more susceptible to the coronavirus, the xinjiang concentration camps are worse than the U.S., therefore the xinjiang concentration camps must have been hit harder by the coronavirus." You could just as easily take the numbers from January. "213 dead in mainland China but 0 dead in Guantanamo? Infinity times as many people dead in an economically prosperous China than in a U.S. prison camp? Those numbers from the U.S. must be fake."

I also fail to see any reason why the Xinjiang concentration camps wouldn't have successfully contained the coronavirus. We know that it's possible to contain the coronavirus based on places like Taiwan and South Korea. We know that China has certainly taken mobility restrictions, mask usage, and test & trace measures very seriously. We also know that the concentration camps, horrible human rights abuses that they are, are at least places where the CCP can enforce drastic restrictions and monitoring on their occupants movement, apparel, and behavior. Why wouldn't they be able to stop the spread in Xinjiang?

None of this is to defend the CCP, or deny that the regime is sketchy, or that the imprisonment of muslims in Xinjiang isn't a horrible abuse of human rights. The point is that the reasoning you've suggested is pretty silly, and doesn't convince me that distrust of Chinese numbers which have thus far been consistent with global information is anything more than a justified distrust.

Joshua on May 8, 2020 1:04 PM at 1:04 pm said:

John –

> But pay close attention to what I did include in the excerpt above:

You included a description of his model output that was highly misleading.

Reply ↓

dhogaza on May 8, 2020 1:33 PM at 1:33 pm said:

"Last March, Ferguson admitted that his Imperial College model of the COVID-19 disease was based on undocumented, 13-year-old computer code that was intended to be used for a feared influenza pandemic, rather than a coronavirus."

I haven't looked to see what kind of model Ferguson has implemented, but the fact that it is 13 years old is irrelevant. It's not like there has been a sudden revolutionary revelation in the mathematics underlying such models in the last 13 years.

The fact that it was developed for an influenza pandemic is almost certainly irrelevant. It would only be relevant if the parameterization of the model was hard-wired into the model, i.e. rate of infection, for instance. Since we know the model was run with an R0 of 2.4, and sThince we know that influenza has an R0 in the neighborhood of around 1.5 or so, obviously the model wasn't hard-wired for the flu. Duh.

"undocumented" you use undocumented or poorly documented, poorly structured code not written to (say) FAA standards every time you touch your keyboard. It is not, in and of itself, and indicator of unreliable code.

The real issue is whether or not the model has made accurate projections.

Here, the model has been disappointing. Only 7,000-20,000 UK deaths in a lockdown scenario? 35K and counting, bro.

Reply ↓

Bob on May 8, 2020 1:43 PM at 1:43 pm said:

Right, so there's only 'undocumented' or 'FAA standards'. Nothing in between.

"I haven't looked to see what kind of model Ferguson has implemented"

Maybe you should. As it stands your opinion is uninformed.

Reply ↓

dhogaza on May 8, 2020 2:52 PM at 2:52 pm said:

Bob:

"Maybe you should. As it stands your opinion is uninformed."

Cool. In order to help me out here, can you describe what the model does?

Thanks.

Tim on May 8, 2020 11:01 AM at 11:01 am said:

Did you see also though, that John Carmack helped tidy it up, and didn't find it so bad? If we're talking about track records, I'll listen to Carmack over 'lockdownsceptics'.
I have to admit my view is it's too easy for academics to write broken code and get away with it, so kudos to Ferguson for releasing a version of it eventually.

Reply ↓

Bob on May 8, 2020 11:43 AM at 11:43 am said:

Hey, I'm not reading that, I might end up agreeing with it. I'll listen to this other guy who says it's nothng to worry about.

You see, there's the problem. You don't have to 'listen' to anyone. Personally, I read the critique point by point and decided if they were valid. There are links provided to follow up if you choose to do so, rather than appealing to authority because you're afraid of the conclusion it might lead you to. Besides, Carmack's defence amounts to 'hey, it's not that bad, I've seen worse'. Let's see if Boeing produce software so terrible and see if anybody is willing to make such a weak defence of it. This is far more important that the usual academic pootling.

Besides, it's not his code, it's the version that Microsoft spent a month cleaning up. It's still bug ridden, to the point you don't get the same result with the same random seed. The difference is 60,000 deaths in one example given.

The Imperial group, rather than fixing this fundamental issue, decide instead to call this 'non-determinism', and average over many runs. This is all in the issues list on github.

If the 'refactoring' has introduced the bug, why doesn't he release his original code? Because it's trash, they don't understand what it does. Not a surprise in academia, but his model is the reason the UK is locked down.

But to the original question, Ferguson got traction because he contradicted the then policy and gave a prediction which looked like a doomsday scenarion. It's what people, i.e. The BBC, Guardian, Financial Times etc. wanted to hear. The Guardian want to belive the Tories want to wipe out poor people and the disabled. The FT think everything is about Brexit, and want to keep fighting that battle. The left wing element of the UK media despise Dominic Cummings, and try to convince everyone he's the next Goebbels.

Ferguson gave them what they wanted, his input 'proved' the government 'wanted thousands of people to die'.

He was treated like an oracle, recieved hagiographic treatment in the FT in particular. It fed onto the narrative that the UK government had made a massive blunder. That picture has become less certain since the actual out-turn of deaths, the redundancy of the surge capacity that was built, and events in Sweden.

A lot more detail has emerged about the scientific advice, which is all minuted and online. It turns out Ferguson was not some lone voice in the wilderness, but was actually part of the scientific advisory process long before his model became public knowledge. It turns out the UK government has been following the advice it was given, so the initial framing of the UK media has been shown to be false.

Accordingly, Ferguson's profile has fallen a lot in the last few weeks. I doubt this is unconnected with the scrutiny his model and in particular his inability or unwillingness to publish his code has received. In fact, he's persona non grata now because he was breaking the lockdown rules he reccomended in order to meet his married girlfriend. He has resigned, but he wasn't useful to the media anymore anyway. They've moved on to sowing confusion around the end of the lockdown.

Reply ↓

Twain on May 8, 2020 12:00 PM at 12:00 pm said:

+1

Reply ↓

Joshua on May 8, 2020 12:06 PM at 12:06 pm said:

> But to the original question, Ferguson got traction because he contradicted the then policy and gave a prediction which looked like a doomsday scenarion. It's what people, i.e. The BBC, Guardian, Financial Times etc. wanted to hear. The Guardian want to belive the Tories want to wipe out poor people and the disabled. The FT think everything is about Brexit, and want to keep fighting that battle. The left wing element of the UK media despise Dominic Cummings, and try to convince everyone he's the next Goebbels.

Ferguson gave them what they wanted, his input 'proved' the government 'wanted thousands of people to die'.

————————–

This is ridiculous spin.

Reply ↓

Bob on May 8, 2020 12:11 PM at 12:11 pm said:

How is it 'ridiculous spin'?

Reply ↓

Joshua on May 8, 2020 12:17 PM at 12:17 pm said:

What is your evidence that determines the direction of causality you outline?

> It's what people, i.e. The BBC, Guardian, Financial Times etc. wanted to hear.

On what basis do you determine that people "wanted to hear" that tens of thousands of people would die?

> The Guardian want to belive the Tories want to wipe out poor people and the disabled.

Where did Ferguson's model output say that Tories wanted to wipe out poor people and the disabled? How did his output "prove" that the government wanted thousands of people to die?

His model output was based on conditional probability – that absent certain interventions there was a range of outcomes – including low end outcomes that would have not been a "doomsday scenario."

You're seeing what you want to see, after you've spun it to match your political predisposition.

Reply ↓

Joshua on May 8, 2020 12:18 PM at 12:18 pm said:

In short – motivation impugning requires a high bar of proof.

You've provided none.

Joshua on May 8, 2020 12:38 PM at 12:38 pm said:

Bob –

I'm going to "prove" to you that rightwingers "want" people who "stay at home" to die.

Remember conditional probability is hard.

In rightwing web comment pages, there are a lot of people making a big deal out of a recent survey that found that recently, 66% of hospitalizations were people who had been “staying at home.”

Sounds really bad, right? It proves that “lockdowns” don’t work, right?

But I don’t see those folks making a big deal out of this asking what % of New Yorkers would describe themsevles as staying at home. Imagine it is 80%. If so, what would it mean that 66% of hospitalizations were among people who self-report “staying at home.”

Further, such a survey of self-report is notoriously likely to be inaccurate.

Now imagine that people were saying that the rightwingers “wanted” results such as this survey because it “proved” that rightwingers want people who “stay at home” to die.

That is essentially the logic that you’re employing above.

Bob on May 8, 2020 12:54 PM at 12:54 pm said:

“On what basis do you determine that people “wanted to hear” that tens of thousands of people would die?”

I do it by reading the comments under the articles in the Guardian and FT for example.
——————————————————————————————————

“Where did Ferguson’s model output say that Tories wanted to wipe out poor people and the disabled? How did his output “prove” that the government wanted thousands of people to die?”

I never said it did. If however, you are a person with strong priors that the tories hate poor people, or hate disabled people, like a lot of the Guardian readership, you may take it as evidence that you are correct. Many in the comment sections did exactly that.

——————————————————————————————————

“His model output was based on conditional probability – that absent certain interventions there was a range of outcomes – including low end outcomes that would have not been a “doomsday scenario.”

His intervention was to say that the then government strategy would lead to a range of outcomes, the high end of which was 500k deaths. Is this not a ‘doomsday scenario’?

——————————————————————————————————

If you want another example, look at the Guardian treatment of the IHME model. From front page splash to being dropped down the memory hole a couple of days later.

Why did it make the front page? Easy, it suited the view of that particular paper.

Why was there no informed critique of this model? After all, there were plenty of people pointing out on twitter that the model was wrong in terms of basic things like the number of ICU beds. Not predicted beds, beds on the actual day of the prediction.

For the next couple of weeks there were people pointing out the gaping holes in that model as they made their updated predictions, there was no Guardian front page revisiting an obviously badly flawed model. Why not?

After all, I saw people on twitter quoting it even then, even after it had become a laughing stock. Surely the Guardian would want to correct any misconceptions it had contributed to?

I know why, and anyone who watches what newspapers publish know why.

Maybe it's news to you that newspapers tend to reflect back to their readership their pre-existing views, but I thought this was common knowledge.

Joshua on May 8, 2020 1:02 PM at 1:02 pm said:

> I do it by reading the comments under the articles in the Guardian and FT for example.

How did you control for your own political biases in your selection and evaluation criteria?

> "Where did Ferguson's model output say that Tories wanted to wipe out poor people and the disabled? How did his output "prove" that the government wanted thousands of people to die?"

>> I never said it did.

You said that people used the model output to see "proof" of such at thing. So you're saying that they used his model to "prove" something that wasn't in his model? If so, how was he doing what they wanted him to do by producing a model that didn't do what they wanted it to do?

> If however, you are a person with strong priors that the tories hate poor people, or hate disabled people, like a lot of the Guardian readership, you may take it as evidence that you are correct. Many in the comment sections did exactly that.

Sure. If your priors are strong enough, you can use his model to prove whatever you want. Many have used his model to prove that libz want to panic everyone so they can take away freedoms.

> His intervention was to say that the then government strategy would lead to a range of outcomes, the high end of which was 500k deaths. Is this not a 'doomsday scenario'?

No, because it can't be decoupled from the low end of the range, which is what you're spinning to do. A "doomsday scenario" would be the 500k number without the low end numbers of equal (or greater) probability.

> If you want another example, look at the Guardian treatment of the IHME model. From front page splash to being dropped down the memory hole a couple of days later.

I'm not going to engage with you doing more of what is the same thing.

> Maybe it's news to you that newspapers tend to reflect back to their readership their pre-existing views, but I thought this was common knowledge.

I completely agree that's what newspapers do. I'm talking about something else – your motivation impugning.

By using your logic, I can "prove" that rightwingers want people to get sick and die under lockdowns so they can prove that lockdowns don't work.

That's the kind of trouble you get into when you impugn motivations without meeting a high bar of evidence. Divining motivations is difficult – particularly when you, yourself, have a strong ideological orientation.

Joshua on May 8, 2020 1:26 PM at 1:26 pm said:

Bob –

I'll leave you with one more comment:

> Maybe it's news to you that newspapers tend to reflect back to their readership their pre-existing views, but I thought this was common knowledge.

Keep in mind that commenters at newspaper websites are outliers. The vast majority of readers never comment, and the commenters tend to cluster at the extreme ends of the ideological spectrum.

Extrapolating from comments to the wider readership is not well founded.

dhogaza on May 8, 2020 1:46 PM at 1:46 pm said:

Bob says:

"If you want another example, look at the Guardian treatment of the IHME model. From front page splash to being dropped down the memory hole a couple of days later.

Why did it make the front page? Easy, it suited the view of that particular paper."

Trump began touting it, too. Apparently he and the Guardian share the same political views.

Help me out here. I'm a yank. Explain to me the shared political views held by the Guardian and Trump.

On May 3rd, the IHME model was predicting about 74,000 deaths in the US by August 4th. I am posting on May 8th. We will end the day at around 79,000 deaths.

The IHME model sucks (or did, they've rewritten it, it's no longer purely an exercise in curve fitting, so who knows, it might perform better going forward). But that suckiness led to its projections being constrained mathematically to a value that would be too low unless deaths would drop drastically once the peak was reached. Regardless of the data. This doesn't fit your narrative that model inaccuracies have led to unwarranted concern and governmental policies.

dhogaza on May 8, 2020 12:24 PM at 12:24 pm said:

Bob:

"Ferguson gave them what they wanted, his input 'proved' the government 'wanted thousands of people to die'."

Huh, 35K have died in the UK. And the number's still growing. If the model code were written to industry standards, no one would be dead?

Reply ↓

Bob on May 8, 2020 12:57 PM at 12:57 pm said:

What are you tring to achieve with that non sequitur?

Seriously, this isn't a difficult concept. The question is why Ferguson got so much traction.

Reply ↓

dhogaza on May 8, 2020 2:09 PM at 2:09 pm said:

'Seriously, this isn't a difficult concept. The question is why Ferguson got so much traction."

Well, for one thing, the UP TO 500K UK/2.2 US deaths are actually about what one would expect, and he and his modeling efforts were not the only source of it.

Back in late February, the WHO published a report on their initial epidemiological estimate of the infected fatality rate for covid-19. This was a very rough estimate, but it has held up quite well.

0.3%-1.0% (for one comparison, serological results in NYC lead to computed IFR in the neighborhood of 0.8%).

Based the estimated R0 value available at the time, epidemiologists were already stating that "herd immunity" would require 40%-60% of the population become infected.

The lowest number of estimated deaths in the US with these parameters (40% infected required for herd immunity, IFR of 0.3%) would be just under 400K.

The highest would be about 2 million.

This is absent any mitigation measures, including those taken voluntarily ("head for the hills!").

So Ferguson's "up to 2.2 million" really is not an unexpected number. Yes, it met the expectations of a crowd of scientists, but for good reasons.

If he'd said "up to 10 million dead in the UK" I doubt very much that he would've gotten traction, because even back-of-the-envelope calculations by others would make it clear that this number was clearly ridiculous.

Oh, yeah, but you can't read the code. That could be taken, you know, as a lack of effort or skill or both on your part.

People have said the same about GISS Model E (same arguments, and with the same motives, i.e. to undercut the credibility assigned to this and other GCMs). Oddly, when I dug into the source, I was able to read it fine. I don't have the physics background to vet it in that way, nor am I qualified to validate the various parameters that simplify things that simply couldn't be modeled given the granularity used to run the model for the entire planet in reasonable time on the supercomputers available to them at the time. And I was last fluent in FORTRAN back in the early 1970s, which didn't help. But the code quality complaints were far, far overblown.

Reply ↓

Ben on May 8, 2020 1:12 PM at 1:12 pm said:

> Let's see if Boeing produce software so terrible and see if anybody is willing to make such a weak defence of it

https://en.wikipedia.org/wiki/Boeing_737_MAX#Grounding_and_suspension_of_production

That was a software thing Boeing did.

Reply ↓

Bob on May 8, 2020 1:38 PM at 1:38 pm said:

Yeah, I know. Now let's see if anyone defends it.

Reply ↓

Ben on May 8, 2020 2:17 PM at 2:17 pm said:

Ohhh, I think we're reading this two different ways. I see this as an outlier where the software was bad enough that it actually triggered an investigation and forced Boeing to take the blame. I think of this as evidence that most stuff out in the wild is just barely hanging together.

I think the way you're reading this is that this is evidence that if there is bad software in the wild, it will get tracked down and fixed.

Reply ↓

Zhou Fang on May 8, 2020 3:09 PM at 3:09 pm said:

The difference is 60,000 deaths in an individual run – but the intended idea of the program is to generate many thousands of iterations of these runs and aggregate to compute intervals.

The bug was that the RNG was being re-initialised at some point in the run and not being saved, which does not affect the theoretical validity of the model at all (though it does make debugging potentially more difficult).

Ferguson's model is not the only one, and it generally agrees with what other experts predict.

Frankly getting out a prediction with roughly the same numbers as Ferguson is not difficult. Assume that the UK ends up with 50% infected (the herd infection limit). Multiply that by population and IFR (let's say 1%). Boom, that's about 300k dead. The key uncertainties are in those quantities, and those have nothing to do with little things like RNGs.

You think the proportion of people who die from covid is way less than 1 in a scenario with millions infected%? You think the progress of the infection will suddenly stop somehow at much less than 50% (like 5%?) without intervention? Those are useful points to discuss. This discussion sounds like the sort of fixation on tiny details to try and bust the entire argument that is typical of conspiracy theorists.

Reply ↓

More Anonymous on May 8, 2020 3:23 PM at 3:23 pm said:

+1

Reply ↓

Joshua on May 8, 2020 3:29 PM at 3:29 pm said:

Zhou –

> Assume that the UK ends up with 50% infected (the herd infection limit). Multiply that by population and IFR (let's say 1%). Boom, that's about 300k dead. The key uncertainties are in those quantities,

Indeed. When I first heard about these model outcomes I pretty much assumed that's what was done. Same for the projections for the US.

As such, I reflexively view the targeting of the IC modeling methodology as spin. We don't really need complicated models to make a basic "if nothing is done to mitigate" projection. Of course, that should be done only with an understanding that the "herd rate" % is, as yet uncertain, and the IFR is (as yet) uncertain. Seems to me that the value of the modeling is to help quantify the implications of the uncertainties – that's not something that is easily done through simple calculations.

Reply ↓

Zhou Fang on May 8, 2020 3:37 PM at 3:37 pm said:

Yes, this was the original context of the issue being raised in the github by the red team. They wanted to conduct systematic parameter sensitivity testing, so removing other sources of uncertainty helps simplify that effort. The issue with the reseeding of RNGs was fixed two days after it was raised but I guess this means the lockdown is permanently unjustified now :P.

Reply ↓

Joshua on May 8, 2020 3:50 PM at 3:50 pm said:

Zhou –

> They wanted to conduct systematic parameter sensitivity testing,

So as someone who's here to learn…

Seems to me that conducting parameter sensitivity testing only makes sense for modeling that is so dependent on highly sensitive and uncertain parameters – which I assume to be the case (as someone not statistically conversant) since the range of the projection is so large.

Also, it seems to to me that critique on those grounds is ENTIRELY legitimate. A more broadscale critique just seems likely to me to be politically motivated spin – because again, we can just use a very simple common sense project # of infected multiplied by IFR.

Then, it seems to me it's time to move on since modeling of how the infections are likely to play out seems like where we should be shifting, as that would have the most direct implications to policy development…

But that would be in a country with competent leadership…

Anyway…

–snip–

This Is the Future of the Pandemic
Covid-19 isn't going away soon. Two recent studies mapped out the possible shapes of its trajectory.

https://www.nytimes.-com/2020/05/08/health/coronavirus-pandemic-curve-scenarios.html

Stephen M St Onge on May 11, 2020 10:24 PM at 10:24 pm said:

Bob:

You touch on, but miss, what I think is the fundamental point: NO VALID PREDICTION WAS POSSIBLE.

Because of the uncertainties in factors such as Ro, IFR (which varies with health, age, and sex), size of the population at risk, probability of exposure, prevention measures, changes in behavior as the epidemic progresses, and the availability of medical care, the best that could be produced was a probability distribution. The news media chose to take the greatest figure produced in the distribution, and label it a "prediction," but it was actually only an estimate of a possible extreme.

I'm reminded of four military 'predictions' that went catastrophically wrong. The British "10-year rule" was the result of asking 'How long will it be before there is a possibility of another major war?' That was asked in 1919, and the answer was 'Ten years from now.' The next day, the next month, the next year, and ultimately the next decade came, and the answer remained 'Ten years from now.' It was about 1931-3 before the Brits woke up and realized they'd blown a decade of possible preparation.

The same thing happened with the U.S. trying to predict when the USSR would acquire nuclear weapons. The answer was always 'X years from now', and it too became a moving target X years in the future. What was needed in both cases was an answer in terms of dates:'Not before 1929, but possibly then, and certainly Germany will be fully recovered by 1932′ for example.

The other two were Soviet response to Hitler's invasion build-up, and the U.S. response to the Japanese obvious preparation for war, both in 1941. The Soviets asked if the Germans _were_ planning to attack, while the U.S. asked _where_ the Japanese would attack. The proper question for the Soviets was 'If those forces do attack, what happens?' (Answer, they get over-run, because there are NO defensive precautions across most of the border). The U.S. should have asked 'What happens if they attack at location X, using method Y?' (Answer for Pearl Harbor: disaster, because Kimmel and Short REFUSED to prepare for an air attack.)

The great mistake of the media and the public is not understanding this, and asking for any prediction at all. By looking for a prediction of what _would_ happen, they failed to consider the various "might happens." Concentrating on those might have led to a plan with intensive monitoring capability, and the capability of quick adaptation to changing circumstances.

Reply ↓

dhogaza on May 8, 2020 12:35 PM at 12:35 pm said:

The whole "code is not written to industry standards" argument is copied from the climate science denialism playbook.

It went something like this:

"free the code!" "we can't trust your models unless you free the code!" (regardless of how well model projections fit reality)

Them when the code (GISS Model E, specifically) was released:

"the code doesn't meet industry standards!" "OMIGOD it is written in FORTRAN! It must suck!" "Hardly any comments, I can't understand it, it must suck!" "I can't get it to compile on my box! It must suck!" (regardless of how well model projections fit reality, because that's the goal is to discredit such models, not determine if they're useful or not).

And normally the standards being discussed are for aerospace applications that have to meet FAA standards, etc.

These people never seem to mention that most of the software industry doesn't write code that meets whatever arbitrary standard these people hold up as being necessary for the production of reliable software.

If you think software needs to meet "industry standards" to be of reliable or useful, I suggest you don't read the linux kernel.

Now, the IC model has had some obvious issues. James Annan (who's posted here a couple of times) raked them over the coals for sticking to an R0 of 2.4 when evidence made clear it was around 3. And quite a few other things. Based on results and model methodology. You don't need to read the code to understand this. And even code meeting the most stringent standards that's built on poor assumptions or modeling methodology is going to be crappy.

Reply ↓

dhogaza on May 8, 2020 12:36 PM at 12:36 pm said:

Pardon the typos above ...

Reply ↓

Rahul on May 8, 2020 12:57 PM at 12:57 pm said:

I somewhat agree with a part of what you write. e.g. the Aviation / Nuclear etc. coding standard is a bit ridiculous for such things.

OTOH, I draw a distinction between epidimological / climate models etc. vs more functional code. i.e. If you write a device driver or kernel or even a Pacman game the proof is in the pudding. If the code crashes or behaves weirdly it is pretty obvious and your users will notice. And the system corrects itself or goes to the graveyard of failed projects.

Models are different. The "fits reality" bit is somewhat difficult to judge. You don't have huge data and are always in some ways extrapolating. So for those it is fundamentally important that your code be readable enough that others can vet it.

Another example of this is cryptographic software etc.

So yes, I hold a model to a different standard.

Reply ↓

Daniel Lakeland on May 8, 2020 1:46 PM at 1:46 pm said:

In the abstract, I agree with you. But in the scenario "You have about 10 days to make a major decision which if you do it wrong could costs literally hundreds of thousands or millions of lives" what do you do?

The main question is more like this: is the ratio of deaths under the different scenarios reasonable? If the UK hadn't locked down would it have had on the order of 10x or 20x or more as many deaths/morbidity cost as we did under lockdown?

It seems pretty clear to me that the answer is yes. People with political reasons for doing so are basically saying "things would have been basically the same even if we didn't lock down" with *no* backing for it, and pointing to "but the code" as their justification why the choice was wrong... without *any* justification why their "gut feeling" about "things would have been the same" should be given *any* weight at all.

Reply ↓

Rahul on May 8, 2020 1:53 PM at 1:53 pm said:

@Daniel:

I would have agreed with you had we lived in a world where normally most academic modeling code was well written, extensively tested, independently reviewed published transparently and all the other nice things. And only because of the "10 day deadline" of Covid we had to cut corners.

But we don't. Academic models are a mess in the best of times. Ferguson is not an exception with the excuse of Covid time pressures.

My point is, this low standard, is par for academic models and code. Covid or not. Covid just brought things into the spotlight.

Reply ↓

Daniel Lakeland on May 8, 2020 2:02 PM at 2:02 pm said:

The time pressure is: make a decision "today"

what will you use for making that decision? The best available thing right? If the best available thing is some not that great code from a decade ago… then that's a lot better than "Boris Johnson, who has no knowledge, says it's probably going to be ok, let's do herd immunity"

If you want academic code during normal times to be held up to some kind of robust standards, set the standards, pay for it to be done, and audit that it is done.

there are lots of examples in biology of big packages of bioinformatics code produced to relatively high standards. The fact is, the grants are there *to produce the code* and the users are there to file the bug reports.

Reply ↓

Rahul on May 8, 2020 2:15 PM at 2:15 pm said:

Re "pay for it to be done". Correct!

Right now, funding agencies are "paying for" lots (too many) of low quality code. They should focus on funding fewer, higher quality projects.

If a model is worth funding to be written, it is worth funding to be written well and vetted.

dhogaza on May 8, 2020 2:19 PM at 2:19 pm said:

The reality is that all we need to know is whether or not the code properly implements the model.

So the first place to look is at the model itself. I assume that Ferguson has published?

You don't need to read the IHME code to realize that their exercise in curve fitting was … dubious … at best. They described their model in a high-level online document and a more detailed paper preprint. Once professionals saw this criticism began flowing freely.

So seriously the code isn't the first place people should be looking. Doesn't matter if the code properly implements the model if the model sucks. The first release of the IHME model was open sourced to github, written in python IIRC (using standard math packages), didn't look horribly written, and that's irrelevant.

So as a first step, the question should be is whether or not the IC model is a reasonable one. At that point, trying to figure out if the code implements that model accurately is important.

But the self-appointed code-reviewing ankle-biters don't even seem to have an interest on whether the model is reasonable or not. Bang! The code! Looky! Again, as I've said above (and Joshua, too) this is taken right out of the climate science denialist playbook.

And to paraphrase Galileo … and yet, people are dying.

More quickly than the model projected in the lockdown scenario …

Daniel Lakeland on **May 8, 2020 3:06 PM at 3:06 pm** said:

dhogaza: I agree with you. In fact, it doesn't even need to be implemented correctly (nothing of this size is ever completely correct), all it needs is that the bugs don't produce qualitatively terrible errors.

qualitatively, all models of the COVID dynamics agree, if you compare "do nothing" to "lockdown" the results will be dramatically fewer deaths under lockdown. And, if you compare "do something kind of like Sweden" to "lockdown" the results will be dramatically fewer deaths under lockdown… and all the data we have is consistent with that.

The only way we can use models to get more refined information than these qualitative ones, is to do full on Bayesian analysis of an ensemble of models, a process that can't be vetted in the public eye because the public doesn't have anything like the expertise… IT *should* be vetted in the public eye of experts though. And IC put up some kind of model on github which is exactly what they should have done. Basically all of this political crap is just that… crap. People want to arrive at the conclusion "open all the things" and they'll do anything they can to justify that by any means possible.

the problem with the IC model is similar to the problem of GCMs they have *too much detail* which makes them attackable via the death of a million cuts. The right level of modeling is sufficiently simple to capture the main important dynamics

confused on **May 8, 2020 5:58 PM at 5:58 pm** said:

>>If the UK hadn't locked down would it have had on the order of 10x or 20x or more as many deaths/morbidity cost as we did under lockdown?

>>It seems pretty clear to me that the answer is yes.

I'm not sure how this can be known with even reasonable assurance at this point.

Sweden isn't doing great, but it actually has slightly fewer deaths / million population than the UK.

The US doesn't show a strong trend between states with strict shelter-in-place / stay-at-home orders vs. those with more moderate social distancing, either.

This could be because the areas that didn't lock down were inherently less vulnerable due to lower density / less social contacts at baseline, or it could be evidence that moderate social distancing is enough to prevent hospitals from being overwhelmed in most places (not Madrid or Lombardy etc., and maybe not NYC).

I don't think the answer is knowable at this point, until some of the denser areas relax their orders.

Reply ↓

Brent Hutto on May 8, 2020 6:09 PM at 6:09 pm said:

It's obvious that other factors are far more important in determining whether hospitals get overwhelmed than just the timing and thoroughness with which lockdowns were implemented.

It appears the full answers to why some places have health system crashes and other don't are not knowable at this point. I think it's much the lack of reliable and detailed data as it not having densely populated areas relaxing their lockdowns.

Reply ↓

Stephen M St Onge on May 11, 2020 11:14 PM at 11:14 pm said:

Daniel Lakeland:

"The main question is more like this: is the ratio of deaths under the different scenarios reasonable?"

BINGO! i have a continuing 'argument' on FailBook about when the Chinese Commie Virus showed up in the U.S. This woman insists it arrived in November, and spread widely before the first reported case in January. It must have happened because of the available flights from China, and people remembering flu-like symptoms in December.

I respond that with a three-day doubling time, if the virus arrived in November, all of us would have had it by the end of February (1,000 to the third equals three U.S. populations), and the antibody testing shows only a few percent have been infected nationwide. She comes up with new anecdotes. She knows what she wants to believe, and ignores the fact the data contradicts her.

All this stuff about Sweden is another example. It's impossible to get quantifiable evidence for social distancing in Sweden (or at least, I can't find it). News stories say the Swedish infections are heavily concentrated in the small immigrant population, which every concedes refused to follow the Swedish governments advice. (https://www.h-stoday.us/subject-matter-areas/counterterrorism/when-religion-and-culture-kill-covid-19-in-the-somali-diaspora-communities-in-sweden/). Sweden's death rate is six times that of Norway, which also has Muslim immigrants, but a smaller proportion of the population. Depending on the assumptions I make, I can 'prove' Sweden's policy is a great success among everyone but the immigrants, or an absolute disaster.

I never realized how strongly people believe things because it fulfills an emotional need, rather than being indicated as true by facts and logic.

Reply ↓

Daniel Lakeland on May 11, 2020 11:57 PM at 11:57 pm said:

> I never realized how strongly people believe things because it fulfills an emotional need, rather than being indicated as true by facts and logic.

The extent of this has been striking to me as well over the years since Trump was elected, but especially in the last few months.

Reply ↓

Joshua on May 12, 2020 12:12 AM at 12:12 am said:

Some anecdotal information about social distancing in Sweden:

https://www.nationalgeographic.com/history/2020/05/sweden-hasnt-locked-down-but-normal-life-is-a-luxury/

Reply ↓

K Yong on May 27, 2020 6:17 PM at 6:17 pm said:

@Daniel

As a total layperson taking interest in the discussion I see the vast complexities of the pros and cons of applying a predictive model based on assumptions and limited data.

What intrigues me is the scenario you brought up:

"You have about 10 days to make a major decision which if you do it wrong could costs literally hundreds of thousands or millions of lives" what do you do?" (It seems like something the Joker would say to Batman to force his hand).

My answer. Get an opposing view that accounts for the consequences that you want to take. Factor in socio-economic costs into that comparative model. Use the median value or predicted deaths to frame the problem instead.

People seem to forget that decision-making, especially in times of crisis, is both a heightened intellectual and emotional exercise. There's never only 2 factors or a simplistic cause and effect.

In my mind, the decision making would have been significantly different if there had been two contrasting predictive models to contemplate. One of which would take into consideration collateral damage (i.e., the socio-economic fallout that would manifest, increased suicides, morbidities from people not seeking medical attention for fear of getting the virus, homelessness, domestic abuse, chronic unemployment, civil unrest, loss of individual rights, the rise of a police state, global tensions and upheaval, global food shortages, increase in global poverty and hunger, etc.) that would occur using, in this example, Ferguson's flawed model.

It's this kind of decision-making in a vacuum that is the real issue. The ethical dilemma raised in some of the posts I've read imply that the tunnel vision in the use of scientific data and models is controversial or dangerous.

The bottom line is: how do we compare the cost of human lives lost to the protracted socio-economic toll experienced by those who survive? As an economist would say, what are the opportunity costs?

I'm no expert, just a curious bystander looking for some clarity and truth.

Reply ↓

Ben on May 8, 2020 4:00 PM at 4:00 pm said:

Hmm, I get that this code sounds super sketchy, but the conclusion seem rather strong:

> On a personal level, I'd go further and suggest that all academic epidemiology be defunded. This sort of work is best done by the insurance sector.

The insurance sector could have coded and published such models themselves. Presumably they have lobbyists and media people who could have spread the results from their models?

stats from their models?

> Instead, the Government is dominated by academics who apparently felt unable to question anything done by a fellow professor. Meanwhile, average citizens like myself are told we should never question "expertise".

I don't know if the team turned the author away before, but like, go comment on the code! It looks like they'd appreciate the help, though I don't really know what's going on there.

Part of the article is how this bug gets found by someone: https://github.com/mrc-ide/covid-sim/issues/116#issuecomment-617304550 . The team is skeptical, but asks for a test which the person provides.

Then the team fixes the bug: https://github.com/mrc-ide/covid-sim/pull/121

Like, that's good right?

Reply ↓

Jonathan (another one) on May 8, 2020 10:13 AM at 10:13 am said:

Weak evidence on the futility of lockdowns:
https://www.medrxiv.org/content/10.1101/2020.05.01.20088260v1.full.pdf
The weakness stems from the high collinearity between the interventions studied, but still…

Reply ↓

Jonathan on May 8, 2020 11:35 AM at 11:35 am said:

Thanks other Jonathan for the link. It's an interesting puzzle. Sweden openly discusses that keeping infections out of care facilities proved harder than they expected. And their experience: same 50-60% of all deaths are from care facilities. Same as MA, same as other places. So my reaction is we should be looking at care facilities as a co-morbidity, meaning as a causal factor and not merely as a place where there are clusters. That data also suggests we should be evaluating the non-care facility risk for older people. I can't do that with the numbers I've seen. As in: MA has test numbers and confirmed case numbers, but no longer breaks out age by decade over 80 – perhaps to hide how many centarians are dying (my count was about 50, which is a gigantic number) – so there's no way to break out the data and, importantly, no way to know if they're testing really old people only when they present as ill. You'd think finding a risk number for random old people not in care homes might be a priority, but I see nothing about it.

Anecdotally, what I see is plenty of social interactions: by closing all but 'essential' businesses, we seem mostly to be 'virtue signaling' rather than actually limiting business to essentials. I've had to go to Home Depot twice in the past few days. It was busier on Thursday in lock down than it normally is on any weekday. Everyone was buying flowers, shopping for tile for remodels. I maybe saw one person getting material for an essential repair. I can't buy a birthday card at the card store but I can buy one at CVS, which has been selling lots of them. To me, there is clearly enough social traffic around 'essential businesses' to keep the number of cases spreading. Example: I've been waiting for reports about the virus in the local Armenian community because their markets are open, tiny and crowded with old people.

That is, we confuse a generalized spreading with localized spreading. That's fine: the original goal was to preserve our medical capacity not to prevent people from being exposed. It seems that over the past 2 months, the idea has shifted – perhaps due to the constant scare stories – from bending the curve to not having new cases. It's a virus. There will be new cases.

Reply ↓

Twain on May 8, 2020 12:19 PM at 12:19 pm said:

Jonathan,

+1 and agreed.

> Thanks other Jonathan for the link. It's an interesting puzzle. Sweden openly discusses that keeping infections out of care facilities proved harder than they expected. And their experience: same 50-60% of all deaths are from care facilities. Same as MA, same as other places. […] You'd think finding a risk number for random old people not in care homes might be a priority, but I see nothing about it.

You may find this data interesting: https://www.kff.org/medicaid/issue-brief/state-reporting-of-cases-and-deaths-due-to-covid-19-in-long-term-care-facilities/

"Our data collection effort finds that in the 23 states that publicly report death data as of April 23, 2020, there have been over 10,000 reported deaths due to COVID-19 in long-term care facilities (including residents and staff), representing 27% of deaths due to COVID-19 in those states. Our data also finds that there have been over 50,000 reported cases, accounting for 11% of coronavirus cases in 29 states. In six states reporting data, deaths in long-term care facilities account for over 50% of all COVID-19 deaths (Delaware, Massachusetts, Oregon, Pennsylvania, Colorado, and Utah; Table 2)." (Add NJ to the list of 6 as of yesterday.)

> Anecdotally, what I see is plenty of social interactions: by closing all but 'essential' businesses, we seem mostly to be 'virtue signaling' rather than actually limiting business to essentials. I've had to go to Home Depot twice in the past few days. It was busier on Thursday in lock down than it normally is on any weekday.

This happened to me as well, one week apart. Both times, Home Depot had a line of ~25 people around the building and was apparently packed inside…so we went to Lowes. Self Checkout had a line going to the back of the store and wrapping around to the front. Even the Lumber Checkout had ~15 people in line. Never seen this before, even during springtime. I've seen similar problems at supermarkets (at least three) for over a month: I like to shop at night to avoid crowds; yet it is still packed even at 7:00pm or 8:00pm!

> That is, we confuse a generalized spreading with localized spreading. That's fine: the original goal was to preserve our medical capacity not to prevent people from being exposed. It seems that over the past 2 months, the idea has shifted – perhaps due to the constant scare stories – from bending the curve to not having new cases. It's a virus. There will be new cases.

Not according to many Governors! We can eliminate it if we all work together.

Reply ↓

Joshua on **May 8, 2020 12:26 PM at 12:26 pm** said:

Jonathan –

> It seems that over the past 2 months, the idea has shifted – perhaps due to the constant scare stories – from bending the curve to not having new cases. It's a virus. There will be new cases.

How do you determine what "it seems that?" Whose idea has shifted, due to "constant scare stories."

I see people wanting a lifting of government mandated social distancing to be based on the kinds of metrics outlined by the CDC. They want a robust infrastructure for testing, tracing, and isolating infected people. No such infrastructure exists. In many places, the opening is taking place in spite of the metrics that say that it should only be done after the metrics show sustained improvement.

Where do you see people saying that mandates should be lifted only if there are no new cases?

Reply ↓

Daniel Lakeland on **May 8, 2020 12:27 PM at 12:27 pm** said:

> virus in the local Armenian community

are you in the pasadena/glendale area?

Here's the thing that I'm finding *completely baffling*. The solution was always always always supposed to be:

1) shutdown to reduce the spread rate as fast as possible so that we didn't have multiple places looking like NYC

2) keep the shutdown long enough that we build up testing capacity, and build up a sufficient body of data to know where the outbreaks were most widespread.

3) Simultaneously build up case tracking capacity by a lot.

4) Once case tracking capacity had built up, and hopefully cases per day had declined, reopen with a lot of case tracking and quarantine so that *cases continue to decline even while open*

And yet… in reality we're acting like the plan was 1,2, cured!

it can't be stated strongly enough that the Trump administration has suppressed information from the CDC, purposefully hampered the states in their efforts, including seizing shipments of PPE they've ordered, and intentionally spun all of this to the point where it's essentially active information warfare between feds and state.

But the plan is (or should be) still: 3 and 4 followed by success.

Reply ↓

dhogaza on May 8, 2020 12:45 PM at 12:45 pm said:

"Once case tracking capacity had built up, and hopefully cases per day had declined, reopen with a lot of case tracking and quarantine so that *cases continue to decline even while open*"

I know a state social worker in the Bay area, and they were pressed into doing case tracking after some training. Their county gave up within a couple of weeks because of the volume, and the case load in the Bay area hasn't actually been that bad (compared to LA, for instance, where confirmed cases and deaths per capita became much higher than up north).

So stating that testing and case tracking must be in place implies ramping up capacity, yes, but also generates an implicit requirement that the level of new cases be greatly reduced.

Here in CA, at least, that's been clear all along. "flattening the curve" was an easy message to get out there in a time where things were moving extremely rapidly (how many people in the US had died as of six weeks ago?) but anyone paying close attention to what public health authorities were saying in more detail knew the goal was to push the curve down to get to a very low new case rate ASAP. Which has proved more difficult than they expected because of the rather lip service adherence to SIP.

Reply ↓

Daniel Lakeland on May 8, 2020 1:50 PM at 1:50 pm said:

Exactly. and this is also why they are ramping up case-contact worker training, and *trying to hire 10,000 to 20,000 new contact tracers in the next few months* A factor of something like 30x the current statewide capacity.

The biggest problem in my opinion is that they haven't gotten far enough along in short enough time. That 20k workers goal should be at least at a few thousand by now.

Reply ↓

Nick Adams on May 8, 2020 1:15 PM at 1:15 pm said:

This is the strategy of the "early lockdown" countries – Australia, NZ, Norway, Singapore, etc.
We will see how it turns out.

Reply ↓

Phil on May 8, 2020 1:20 PM at 1:20 pm said:

I agree, Daniel. I hope there's a lot going on behind the scenes, tons of progress being made in several areas, because if instead we are just waiting for some future date when we can go back to normal, well, that's ridiculous.

Reply ↓

Daniel Lakeland on May 8, 2020 1:52 PM at 1:52 pm said:

Me too Phil. I am gratified to hear Newsom give numbers like this in his report from monday.

However, I really think the messaging should be much much clearer, and should have been *right from the start*...

Reply ↓

Twain on May 8, 2020 1:46 PM at 1:46 pm said:

Daniel and others,

Something I saw recently that may be important regarding the mass testing: https://www.medrxiv.org/content/10.1101/2020.04.26.20080911v1?fbclid=IwAR1myUTRDFk-iXsQF_I3rbcsSeTlKRycKPdX1f1P_FPihGc7K0S80IdOztXM

If this article is correct (interested your thoughts about this), then mass PCR-based testing may be problematic.

My personal experience lends me to agree with this article: PCR is very, very sensitive to user-error — which are may increase as we rush current technicians to increase their output and rush new, inexperienced technicians in to perform many tests.

Although, I could not find public EQAs for COVID-specific testing (for PCR, EQAs are crucial because user-error is the main source of false-positives), so there is no way to know for sure as of now.

Reply ↓

Daniel Lakeland on May 8, 2020 1:58 PM at 1:58 pm said:

False positive rates on the order of 2% and due to operator error were the ones my wife gave me (she runs a bio lab, they do PCR all the time). Those are consistent with the IQR of 0.8-4% given in that article.

Such rates are not a serious problem for control of the epidemic. The false positives along with the true ones would be told to quarantine… the epidemic would subside, and a few extra people would be harmed by having to isolate, but if we are doing the isolation right they'd be tested again, and they'd be given govt support (money) to pay for the public good of not spreading the virus. The infection would die down to low levels. S.Korea shows this works.

What's *much more* problematic is high false negative rates. Then, people *with the virus* go out and infect others even though we're doing a lot of testing and tracing. the infection doesn't die down, and all the effort is for naught and we wind up back in lockdowns. Fortunately, again S.Korea shows this isn't a major issue.

Reply ↓

Twain on May 8, 2020 2:15 PM at 2:15 pm said:

Daniel,

> Such rates are not a serious problem for control of the epidemic.

I disagree. But my prior is "mass testing" means "test everyone regardless of major symptoms". If this occurs, but most of the population is not infected, then FPR will cause a much higher count of positives than reality; similar to the Santa Clara Seroprevalence report. This could [1] cause many (not a few) to quarantine or re-shutdown areas that do not need to in reality OR [2] make us think the disease is more prevalent/infection than it really is.

If "mass testing" means "test ONLY those with symptoms", then the above is no longer an issue.

> What's *much more* problematic is high false negative rates.

Agreed. The estimated 1-3d shedding of virus pre-symptoms is also an issue if true; if testing/tracing miss one of a few people, then many cases could result from one error.

Reply ↓

Daniel Lakeland on **May 8, 2020 2:55 PM at 2:55 pm** said:

My point wasn't that there wouldn't be knock-on economic costs to quarantining people without the virus, of course there would be. But it wouldn't lead to you being *unable to control* the virus.

Whereas significant false-negative rates lead to inability to control.

Reply ↓

Twain on **May 8, 2020 3:01 PM at 3:01 pm** said:

Agreed and thanks for clarifying.

Joshua on **May 8, 2020 2:23 PM at 2:23 pm** said:

Twain –

> I disagree. But my prior is "mass testing" means "test everyone regardless of major symptoms". If this occurs, but most of the population is not infected, then FPR will cause a much higher count of positives than reality; similar to the Santa Clara Seroprevalence report. This could [1] cause many (not a few) to quarantine or re-shutdown areas that do not need to in reality OR [2] make us think the disease is more prevalent/infection than it really is.

Must be weighed against the vastly more true positives who would isolate and be less likely to infect others. Don't let perfect be the enemy of good.

Reply ↓

Twain on **May 8, 2020 3:03 PM at 3:03 pm** said:

Joshua,

I don't disagree in principle and wasn't suggesting no mass testing because the FPR is high.

More so it is something to be aware of long-term when making the decisions, since it *could* have substantial consequences depending on conditions.

TBH, though, FPR of PCR-based testing does not matter so much if we have a decent measure of infections via seroprevalence.

Reply ↓

Joshua on **May 8, 2020 3:23 PM at 3:23 pm** said:

Twain –

I only felt compelled to make that point because I read/hear a lot of rightwingers saying that mass testing is useless because it's only a snapshot in time, and that someone who tests negative could test positive a minute later.

Sure – that's true. But the argument – that because it's a snapshot means widescale testing is a waste – is IMO facile because it ignores potential benefit that comes from isolating people who test positive, and contact tracing to identify those whom they may have infected.

My impression is that sometimes you dip your toe into some rightwing fairy tales, perhaps to test them out a bit, and so I wanted to make sure you're considering the other side of the argument that mass testing would mean a lot of false positives walking around.

OTOH, sure – the false positives walking around could have substantial consequences. I'm here to help out with the rightwing fairy tales whenever you need it. :-)

Twain on May 8, 2020 7:23 PM at 7:23 pm said:

Joshua,

How did you know I wear a blonde toupe, tan, and have fairy wings? Seems like you've blown my cover…

Jokes aside, I don't mean to seem like a right-winger! Far from it, honestly. I try to avoid identifying with any side; to me, it just seems to detract from thoughtful discourse. I want to learn and be a thoughtful citizen, that's about it. (Of course, that's what a Republican Operative would say… ;-))

confused on May 8, 2020 7:42 PM at 7:42 pm said:

A lot of people didn't hear more than "flatten the curve to keep hospitals from being overwhelmed".

And hospitals aren't overwhelmed anywhere in the US right now, and it was never close except in the New York area (and maybe Detroit and New Orleans).

Many states (including mine) have not seen really severe outcomes.

So people are thinking the problem is either basically over, or was only ever a huge problem in a few crowded cities, but isn't relevant to where they live.

(Also, a lot of people are saying that there are enough asymptomatic and very-mild cases of this that the vast majority of people who get the virus won't seek testing; therefore even very broadly available testing and contact tracing won't work.

And/or that this pandemic's effect, unmitigated, won't be severe enough to justify the measures required, even if they would work – all you really need to do is protect nursing homes etc., and ask the most vulnerable people outside nursing homes to stay home. (There's an article going around the Internet titled "Woodstock Happened During a Pandemic" that says we didn't respond this way during the 1968 flu pandemic.)

I can't judge whether either of these are true or not. But if you are asking why people are responding the way they are, as opposed to what the *right* response would be, I think it is relevant.

Reply ↓

statflash on May 8, 2020 9:27 PM at 9:27 pm said:

Daniel, about the thing you find baffling — I totally agree. I had been misled to believe that #1-#4 were the game plan. Alas, we are seeing #1, #2, "Mission Accomplished!" (and even #2 is only half-assed). I don't think we'll see a bunch of case trackers get hired. That's another pipe dream. Sad to say, I don't think #3 and #4 will ever happen and at this point it's too late anyway. So I disagree that #3 and #4 should still be the game plan.

Is this not the reality?

https://www.statnews.com/2020/05/01/three-potential-futures-for-covid-19/

Does anyone disagree with this (bleak) outlook?

IMO, I think we're clearly in this reality of 1k-2k deaths per day for 12-24 months. It's going to be massive death. We'll see a lot of our elderly die a year sooner than they would have without SARS-CoV-2 in the world. But it is what it is. Is anyone seriously thinking there can be any containment in the US at this point? Maybe if Democrats were in charge right at this moment, but they are not.

confused – I think you're right about what "a lot of people are saying". I agree that locking down nursing homes (and I'd add halting mass gatherings) is the most important thing at this point. But everyone over 65 needs to come to terms with the fact that they're about twice as likely to die every year now until there is a vaccine that works.

Reply ↓

> Twain on May 8, 2020 9:54 PM at 9:54 pm said:
>
> statflash,
>
> > Does anyone disagree with this (bleak) outlook?
>
> That article raises good points, but biology can be surprising. Remember, similar claims about H1N1 and H3N2. Yet they have not returned to pandemic proportions. Of course, an influenza virus is different from a coronavirus (vaccines, etc.), so it's not a exact comparison per se.
>
> But is possible that SARS-CoV-2 will come and then go like these and other viruses have. Not something to base policy whatsoever (!), but maybe help quell some anxiety about the long-term future for humanity.
>
> Reply ↓
>
> > statflash on May 8, 2020 10:29 PM at 10:29 pm said:
> >
> > Twain, thanks for the modicum of hope. I suppose you're right.
> >
> > Re the Imperial College Model and its prediction of 2.2m deaths … If we do see ~1,500 COVID deaths per day for the next 2 years, that's ~1.1m deaths in the next 2 years. Then Ferguson's model was still too high, but not orders of magnitude too high (as many critiqued) and we did flatten it and spread out the pain. It's pretty clear that the IFR for the age 80+ folks is about 10%. If 60% of all 13m Americans who are 80yo+ had gotten (will get, since it is still wave #1) infected in the first wave and 10% IFR, then that's 780k deaths right there. I don't think it's a stretch of the imagination to say that could have happened in the absence of any interventions!
> >
> > Also, the UK is not doing great at the moment:
> >
> > https://www.theguardian.-com/world/2020/may/04/uk-behind-the-curve-in-curbing-covid-19-deaths
> >
> > so Ferguson being dire was maybe not such a bad thing. Andrew's right, someone needed to play the role of being dire! The world needs Judases and Jesuses, yin and yang.
>
> > confused on May 8, 2020 10:51 PM at 10:51 pm said:
> >
> > I don't think there's any reason for concern about "the long-term future for humanity". For one thing, these things do wane over time even without modern medical interventions.

But even more fundamentally, every civilization in history up until 50-60 years ago had a pretty high underlying risk level from infectious disease (and much of the world still does). We used to have yellow fever and malaria in the US; smallpox and polio used to be common.

(I think this is one reason why we didn't see a lot of dramatic action taken in the 1957 or 1968 pandemics. Everyone "in power" then grew up before antibiotics, before the polio vaccine, etc.)

statflash on May 8, 2020 11:24 PM at 11:24 pm said:

confused — i really appreciate all your points and find them to be enlightening while challenging me a bit in my outlook (the best mix).

>But even more fundamentally, every civilization in history up until 50-60 years ago had a pretty high underlying risk level from infectious disease (and much of the world still does). We used to have yellow fever and malaria in the US; smallpox and polio used to be common.
>(I think this is one reason why we didn't see a lot of dramatic action taken in the 1957 or 1968 pandemics. Everyone "in power" then grew up before antibiotics, before the polio vaccine, etc.)

very interesting point and the historical differences with 50 and 100 years ago is just incredible to think about. "21st century problems and mostly first world problems"

confused, I also agree with you on the plateau of ~1500 deaths per day not being stable. I was assuming an average with a wide variance, ranging say from 100 – 3000 a day. it will be interesting to see if the seasonal change helps as we move into summer. I do predict a decent treatment will probably come before a vaccine. Good point about vaccinations being made available to the elderly first and foremost and to those younger people with underlying immune conditions, etc. Everyone else can wait for widespread manufacturing.

Daniel Lakeland on May 8, 2020 9:57 PM at 9:57 pm said:

Gavin Newsom on monday explicitly claimed to be training 3000 people a week and hiring up to 20k tracers: https://calmatters.org/health/2020/05/california-coronavirus-contact-tracing-training/

Reply ↓

statflash on May 8, 2020 10:10 PM at 10:10 pm said:

It might yet work in California. I have a lot of friends there so I hope it does. But the East Coast?

confused on May 8, 2020 10:34 PM at 10:34 pm said:

I can't see the current plateau of deaths lasting 12-24 months, or really even 3 months. The current plateau is largely an artifact of some states that

plateau is largely an artifact of some states that were hard hit early declining, while others that were barely hit early on increasing (though some states do seem to have a "true" plateau – e.g. Texas, which has had almost the same death rate for a month). It's also distorted by those hard-hit-early states going back and reporting nursing home and other deaths that weren't reported early on when all the focus was on hospitals.

There may be a resurgence as social distancing measures are relaxed, there may be a fall second wave, etc. – but the current plateau doesn't seem "stable" as it is the result of overlapping different states/regions which have very different individual trends.

–

The three possibilities in that article are possible, but they aren't the only possibilities. There are a lot of unknowns.

– Is there a seasonal effect, and if so, how strong is it?

– How much does the R or R0 vary between places with different kinds of social patterns (already differing in the absence of social distancing measures, e.g. NYC vs. South Dakota)?

–

I think better national leadership would likely have improved things somewhat, but not enough to get us to a South Korea/Iceland/New Zealand situation where the virus could be eradicated. All the countries that seem to be successfully following that path have a much greater geographical isolation, a much smaller population, and a culture which is not as suspicious of government control. It might have been 50% better but not orders of magnitude better.

Though Hawaii and Alaska may be able to follow the Iceland/New Zealand path.

–

>>about twice as likely to die every year now until there is a vaccine that works.
Well, until the pandemic ends or treatments become available that significantly reduce the risk of death (which may happen well before a vaccine). Manufacturing issues will still be significant, but it'll be easier to produce enough medicine to treat the fraction of the population at high-risk who actually get infected, vs. enough vaccine to protect the entire population.

Vaccination is one scenario for the end of the pandemic, but other pandemics have ended without a vaccine becoming available, e.g. 1918-19. It's possible that a vaccine will become available only when the pandemic is naturally ending anyway (this is more or less what happened in 2009-10 … though maybe we would have had a third wave without a vaccine, who knows?)

Reply ↓

> dhogaza on May 8, 2020 10:47 PM at 10:47 pm said:
>
> "There may be a resurgence as social distancing measures are relaxed"
>
> Florida new cases have been ticking up for a few days now …

> confused on May 8, 2020 10:56 PM at 10:56 pm said:
>
> Have testing numbers increased?

Here in Texas, new cases have been higher than average over the last few days, but not higher than seen before (4/10 is still the highest day for reported new cases), and the increase in new cases is smaller than the increase in new tests.

In most of the US real new infections have probably been over an order of magnitude higher than those actually detected by testing, so we should expect to see more cases with more testing, even if real new infections are flat (or even if they were declining somewhat!)

It is also probably too early to see much effect from relaxation of measures. If the average incubation period is 4 or 5 days, then add a day or so to get tested, plus time for the results to come back, plus time for the results to end up on the state department of health website…

Effects from people taking the measures less seriously before formal relaxation of measures could be showing up now, though.

Joshua on **May 8, 2020 11:05 PM at 11:05 pm** said:

Deaths and hospitalizations are the more salient measurents as they aren't a function of the level of testing. We will see in one to two weeks what those numbers look like.

Anoneuoid on **May 8, 2020 11:13 PM at 11:13 pm** said:

> *"There may be a resurgence as social distancing measures are relaxed"*
>
> *Florida new cases have been ticking up for a few days now …*

Not according to covidintracking.com which seems to be the best source for US data:

https://i.ibb.co/RHBwMMK/FL58.png

Please start including sources for your info.

Anoneuoid on **May 9, 2020 11:50 AM at 11:50 am** said:

> *Deaths and hospitalizations are the more salient measurements as they aren't a function of the level of testing.*

What are you basing this on? It looks like deaths/day increase at about 0.005 the rate of testing both over time nationally and across the states geographically (pages 3-4):
https://www.doc-droid.net/8ma5TMk/covidstates-pdf

Also, just logically if no one was getting tested then the deaths would be attributed to pneumonia (or whatever) instead of covid so obviously the rate of testing will affect the death counts.

Joshua on May 9, 2020 12:04 PM at 12:04 pm said:

Anoneuoid –

> What are you basing this on?

Um. My advanced degree in armchair epidemiology, of course!

In retrospect I'm embarrassed now by my authoritarian tone.

It was my assumption that cases is a flawed metric in that it is largely a function of the amount of testing and that deaths and hospitalizations would not be so flawed. As you pointed out, I should be more careful about making assumptions.

dhogaza on May 9, 2020 12:06 PM at 12:06 pm said:

Worldometers was my source. Generally thought to be accurate.

However I just plotted the daily confirmed cases data from John Hopkins, and it more closely matches your source.

Joshua on May 8, 2020 10:47 PM at 10:47 pm said:

statflash –

> I agree that locking down nursing homes (and I'd add halting mass gatherings) is the most important thing at this point. But everyone over 65 needs to come to terms with the fact that they're about twice as likely to die every year now until there is a vaccine that works.

Are you underestimating the costs? There are many over 65 who don't live in LTCFs. And many under 65 who have co-morbidities. And minorities who will die at higher death rates. There will be enormous economic costs of all the hospitalizations.

Reply ↓

statflash on May 8, 2020 11:40 PM at 11:40 pm said:

I'm not saying the costs won't be high. There will be excess deaths in the next year or two. I'm just saying I don't see any realistic other future at this point. We've passed the point of no return toward (herd immunity + reasonable public policy on mass gatherings and LTCFs + contact-tracing and good testing in some specific places – maybe HI, AL, CA, OR, American Samoa).

I am no fan of Trump at all, but I think the administration's blundering ineptitude sort of weirdly ended up working as a Taoist "wu wei" approach that is a sad but organic/natural/apathetic "solution" to the pandemic. I agree with 'confused' when he/she wrote above: "I think better national leadership would likely have improved things somewhat, but not enough to get us to a South Korea/Iceland/New Zealand situation where the virus could be eradicated. All the countries that seem to be success-fully following that path have a much greater geographical isolation, a much smaller population, and a culture which is not as suspicious of government control. It might have been 50% better but not orders of magnitude better." Let's admit that America was not destined (not in the right place politically) to lead the world in this pandemic. If 2016 election had been different, but that's just fantasizing about what ifs.

Joshua on May 8, 2020 11:55 PM at 11:55 pm said:

staflash –

I don't think the situation is static. If results in the US prove to be obviously worse than elsewhere, it will have an ongoing dramatic economic effect. This country has enormous resources to bear.

Like with climate change, when the risk remains abstract it is hard to visualize and its easy to diminish the risk for political purposes. If the risk becomes more unavoidable in the redder areas, there will be political will to create a robust infrastructure for testing/tracing/isolating. That's what I meant about underestimating the cost. Right now it's easy to think that opening is basically a no-cost option unless you live somewhere like NYC, or are over 75 with comorbidities. It may well be that in the not too distant future those who seek to diminish the cost of COVID for political gain will find that tactic more difficult to pull off.

Daniel Lakeland on May 8, 2020 11:56 PM at 11:56 pm said:

statflash, I suspect we're going to see some pretty serious situations right soon now in some states. Looking at my graphs of cases per day. It looks pretty bad in:

MN, MD, KS, IL, IN, AZ, NE, WI

All of those are growing super-linearly, though not quite exponentially (total cases are going up either quadratically or maybe cubically at the moment). If they take any kind of letting up path, we'll see transition to exponential growth, and will probably see hospitals overwhelmed in at least some of those states.

Basically things are bad, the best we can hope for in my opinion is some dramatic improvements in care, or possibly some seasonality to the virus, or a mutation that makes it more contagious but dramatically less virulent.

Joshua on May 8, 2020 11:57 PM at 11:57 pm said:

statflash –

> I'm not saying the costs won't be high. There will be excess deaths in the next year or two

And my point is that the costs may well be much greater (and less abstract) than just excess deaths (which is a pretty abstract concept).

statflash on May 9, 2020 12:08 AM at 12:08 am said:

Well said. I can't argue with that. I mean if the facts on the ground do get so shockingly horrific in, say, Indiana, then there could arise some political will "to create a robust infrastructure for testing/tracing/isolating".

As for "those who seek to diminish the cost of COVID for political gain", they do exist for sure, as well as those ___ for economic gain. But what we've seen too is the potential that there exist "those who would believe in diminished cost of COVID, even to the point of their loved ones dying from it, to be used as political pawns for cherished ideological/religious goals". It could go either way in Indiana is what I'm saying.

confused on May 9, 2020 12:15 AM at 12:15 am said:

>> Let's admit that America was not destined (not in the right place politically) to lead the world in this pandemic. If 2016 election had been different, but that's just fantasizing about what ifs.

IMO, it would have taken far more than a different President in office for the US's outcomes to change radically. The really critical mistakes in January through early March were structural and bureaucratic (CDC issues with testing + a bit of "not invented here", and FDA slow-rolling approval of private tests). No President we could have plausibly gotten in 2016 would have fixed these issues; even if they are fixed after this pandemic, I don't think there was a realistic path to get there without a pandemic that seriously impacted the US (2009-10 swine flu was not really seen as that bad).

This is important, IMO, because too many people will think the problems are already fixed if we get a different President this election. But the speed-of-response issues are fundamental to the system, regardless of who is in charge, and would require a fairly thorough redesign of how things work, not just a few dictates from the top.

statflash on May 9, 2020 12:21 AM at 12:21 am said:

Well said, Joshua. I can't argue with that. Thanks for the update on those states, Daniel.

I would just say that even if there is some "political will to create a robust infrastructure for testing/tracing/isolating" due to bad stuff, body bags, undeniable facts on the ground, there will still be the post-truth bs going on. "Those who seek to diminish the cost of COVID for political gain" (and I would add economic gain too) — yes they do exist and are rampant. But don't forget the obverse to that. There are also "those who would buy into the diminished cost of COVID, and I propose even if they see their loved ones dying, to be used for political/economic gains in order to further their ideological/religious agenda". Sadly, they will literally become "diehards". If things get bad in, say Indiana, it could go either way. I don't want it to get ugly, but I'm afraid it might. I hope reasonable people will prevail.

Joshua on May 9, 2020 12:23 AM at 12:23 am said:

confused –

> IMO, it would have taken far more than a different President in office for the US's outcomes to change radically.

It's impossible to say how much better off we'd be if we had a different president. I'm the other hand, it's likely we would have been better off if we didn't have a president who lied and said "anyone who wants a test can get a test" two months ago and continues to lie and say, daily, that we have the best testing in the world (we rank about 43rd in per capita testing even as we rank about 14th in per capita cases).

Yes, there would have been problems regardless, but a president can't correct for problems if he constantly lies and says the problems don't exist, and his political base doesn't hold h accountable for his failures.

statflash on May 9, 2020 12:26 AM at 12:26 am said:

confused — good point about the structural issues being much bigger than just the President or any one administration. We are lucky this pandemic is not as horrible as it could have been and we hopefully we'll get a chance to make the right choices for the next one that will inevitably come at some point.

Joshua on May 9, 2020 12:27 AM at 12:27 am said:

statflash –

> , there will still be the post-truth bs going on.

Yes I agree. And I agree with the rest of your post. Yes, ideologically- and tribally-based motivated reasoning is remarkably powerful.

I guess I was just trying to convince myself by expressing some optimism. 🤔

statflash on May 9, 2020 12:37 AM at 12:37 am said:

Joshua, no shame in being optimistic. I try to find silver linings too in other ways. You are totally right to be upset about the lies about testing from the President. The testing thing was a huge failure as Fauci said as well.

confused, your post below this one: "I am very skeptical that eradicating the virus is possible in the mainland US. It might be *technically* possible, but not culturally or politically. That might change if things got *really* bad in the US-excluding-big-Northeastern-and-Midwestern-cities, but I don't think there is much chance of it getting *that* bad (and if it does, we will probably be close enough to herd immunity that no measures will make much difference to the outcome)."

I appreciate the point of it getting *that* bad being about near the same point as reaching a good level of herd immunity.

confused on May 9, 2020 12:56 AM at 12:56 am said:

I'm not saying we wouldn't have been better off with a different President – I do think we would have been.

But the difference it would have made would have been far, far smaller than the difference made by a more flexible, fast-reacting CDC, FDA, etc.

We will have a new President sooner or later (certainly in 5 years, if not next year). But structural problems don't go away with just time.

Joshua on May 9, 2020 1:05 AM at 1:05 am said:

confused –

> I'm not saying we wouldn't have been better off with a different President – I do think we would have been.

But the difference it would have made would have been far, far smaller than the difference made by a more flexible, fast-reacting CDC, FDA, etc.

—-

Sure. But it's also not entirely clear how separable those two factors are. For example, we don't know how much Trump was an interference, particularly given his attitude towards scientific expertise and the "deep state." Had those institions been given grater leadership, for example, more resources might have been devoted towards creating testing and providing PPE.

Anoneuoid on May 9, 2020 1:18 AM at 1:18 am said:

Until covid-19 hit the US everyone said it was the usual viral pneumonia-induced ARDS, then all of a sudden within a few weeks of it being in the US doctors are defying the hospital administration and turning to social media to say it mimics high altitude illness and blood clots play a huge role in the pathology.

Why did this info only come out once the US was dealing with it?

Anonueoid on May 9, 2020 1:43 AM at 1:43 am said:

And just read these reports about HBOT. Patients **immediately** improve and are asking to stay in the chamber, etc: https://www.hbotnews.org/

If the same BS pushed by the media about ventilators was done for HBOT chambers starting tomorrow this illness would no longer be scary in about a month.

Joshua L Brooks on May 9, 2020 7:33 AM at 7:33 am said:

confused –

Related to what leadership on testing could accomplish

Joshua on May 9, 2020 7:34 AM at 7:34 am said:

https://www.theatlantic.com/science/archive/2020/05/theres-only-one-way-out-of-this-mess/611431/

statflash on May 9, 2020 11:14 AM at 11:14 am said:

About the politics of the mortality rate, from NYT today:

Fight Over Virus's Death Toll Opens Grim New Front in Election Battle
https://nyti.ms/2Lg6U2h

Joshua on May 8, 2020 11:33 PM at 11:33 pm said:

> So people are thinking the problem is either basically over, or was only ever a huge problem in a few crowded cities, but isn't relevant to where they live.

Over time they may change their view on whether it's over. I generally question the value of cross-country comparisons, but this may be worth an exception.

https://twitter.com/DerrickVanGenn2/status/1258455696405868545?s=20

Reply ↓

Joshua on May 8, 2020 11:35 PM at 11:35 pm said:

With the important caveat about number of tests vs. number of cases.

Reply ↓

confused on May 9, 2020 12:20 AM at 12:20 am said:

Yes, *if* it "blows up" in rural parts of the US, views will change. (I do not think it is guaranteed to do so, because some states never put stay-at-home orders in place, just more limited measures, and I think enough time has passed since mid-March to see the outcomes.)

I was describing what I see here right now, not the possible future situation.

I am very skeptical that eradicating the virus is possible in the mainland US. It might be *technically* possible, but not culturally or politically

cally possible, but not culturally or politically. That might change if things got *really* bad in the US-excluding-big-Northeastern-and-Midwestern-cities, but I don't think there is much chance of it getting *that* bad (and if it does, we will probably be close enough to herd immunity that no measures will make much difference to the outcome).

Reply ↓

Joshua on May 8, 2020 4:02 PM at 4:02 pm said:

Jonathan –

> It seems that over the past 2 months, the idea has shifted – perhaps due to the constant scare stories – from bending the curve to not having new cases. It's a virus. There will be new cases.

Again, this comment sticks in my craw (what is a craw, anyway?)

Obviously, it's a completely subjective assessment, and so I can't say that you're "wrong." And I don't know what "it seems that" really means… But I think that the following article makes it clear that at least with respect to "mainstream" of epidemiologists, and major press outlets like the NYT, your statement (as a broad characterization) is just a very poor characterization.

https://www.nytimes.com/2020/05/08/health/coronavirus-pandemic-curve-scenarios.html

They are spending a lot of time thinking about and modeling the very opposite of what you characterized what "it seems" like. They are, in fact, modeling ongoing patterns of new cases.

I'd love to know how you have gotten your impression of what "it seems" like.

Reply ↓

Anonymous on May 8, 2020 5:02 PM at 5:02 pm said:

"…I've had to go to Home Depot twice in the past few days. It was busier on Thursday in lock down than it normally is on any weekday…"

I've been thinking the same way, since the mandatory SIP started. There was a lag b/w the order and distancing at first and finally masks, here in CA. I could bet a small sum (up to 7 bucks) that Costco stampedes and long lines without distance just before SIP orders and about a month afterwards, were more appealing to the virus than a petri dish.

That initial period, before people modified their behavior, was the most dangerous, as it created non-uniform spread. Overall, it's a no-brainer that SIP helped, but as with anything in statistics, it's the distribution that matters, not some average number.

Reply ↓

Twain on May 9, 2020 12:53 AM at 12:53 am said:

> Overall, it's a no-brainer that SIP helped, but as with anything in statistics, it's the distribution that matters, not some average number.

+1. Something I've been wondering — how well did the SIP really work as implemented (not how they *should* have happened)?

NJ, NY, NYC, and DE (my locales) all saw lag in requiring masks, enforcing distancing, etc., after Social Distancing and SIP mandates occurred — some states, like NJ and DE, took ~2-3 weeks to really enforce Social Distancing and masks in supermarkets and big-box stores. Family, friends, and I saw this and were baffled — we thought and discussed "if we aren't enforcing the key aspects of these mandates, how effective will they be"? Further, NYC announced its SIP on a Friday but didn't enforce until a Monday — so there was a rush to buy and horde that weekend, at least from anecdotes I've heard. NJ/DE did similar, if I remember right.

I'm all for Social Distancing, SIP, and other NPIs if they work well. But I also want to know what "work well" means in quantitative terms. That allows us to compare them and their costs to other measures like masks, improved hygiene, etc.

Reply ↓

Joshua on May 9, 2020 12:58 AM at 12:58 am said:

Twain –

> I'm all for Social Distancing, SIP, and other NPIs if they work well.

What if their efficacy can't be measured with a high degree of confidence either way?

Keep in mind, that measuring their differential negative economic impact can't be measured with high confiexme either.

Reply ↓

Twain on May 9, 2020 1:14 AM at 1:14 am said:

Joshua,

I think if implemented consistently, then it's possible to evaluate SIP and Social Distancing to some degree. Now how precise that degree is TBD; we don't know until we try.

But since there was much heterogeneity across the US, perhaps I'm asking for something that isn't possible to quantify well.

And good point — quantifying the true negative-versus-positive impact of SIPs is not practical (IMO). So we should try to find other criteria to use, if possible — like just focusing on infections, deaths, and other measurable non-economic outcomes.

Reply ↓

Joshua on May 9, 2020 7:54 AM at 7:54 am said:

Twain –

I think you'll find these interesting (for the first time ever, I'm finding Twitter to not be a 10th level of hell):

https://twitter.com/mlipsitch/status/1258827506930667523?s=20

https://twitter.com/bansallab/status/1258872610210877442

https://mobile.twitter.com/CT_Bergstrom/status/1256828515942449153

Reply ↓

Twain on May 9, 2020 9:54 AM at 9:54 am said:

Joshua,

Thank you for sharing! Truly strange times when Twitter becomes a useful resource for decent information.

Regarding the first paper:

> "Achieving accurate estimates of heterogeneity for SARS-CoV-2 is therefore of paramount importance in controlling the COVID-19 pandemic."

+1.

"Susceptibility" is an abstract concept that exists for modeling purposes. (My priors are: [a] "a susceptible individual (sometimes known simply as a susceptible) is a member of a population who is at risk of becoming infected by a disease"; [b] "infection" means developing appreciable AND being able to shed enough virus to infect others.) In terms of reality…its much more tricky. There are numerous factors — innate immunity being the major one — that could have major impacts susceptibility and its consequences.

For example, say innate immunity to SARS-CoV-2 is very high among those age 0-40. Given how innate immunity works, it is very possible that their "infection" is so short and so small that they cannot infect others. This would have tremendous consequence on what percent of infection we need for herd immunity AND how transmissible SARS-CoV-2 is overall.

So it is nice to see a paper highlighting how crucial varying susceptibility is and calling for further research on the topic.

Reply ↓

Joshua on May 9, 2020 12:41 PM at 12:41 pm said:

Twain –

Here's a sample for considering IFR:

https://www.usatoday.com/story/news/investigations/2020/05/06/meatpacking-industry-hits-grim-milestone-10-000-coronavirus-cases/5176342002/

10,000 cases, 45 deaths. And we could probably guess more deaths to come from that same pool of cases.

Certainly non-representative w/r/t infection rate. My guess is it might be fairly representative in SES, although not likely so in ethnicity/race. Extrapolating from that maybe no worse than doing what they did with the Santa Clara study – which was to take a non-representative (and non-random) sample and from that calculate an infection rate and then from that extrapolate to a broadly applicable IFR?

But what would the likely age stratification be? My guess is very few over 60, and no one in a LTCF.

confused on May 9, 2020 8:39 PM at 8:39 pm said:

I would say it's almost certainly unrepresentative in terms of SES – meatpacking workers are going to basically by definition be "working class". There aren't going to be any long-term unemployed and/or homeless people in that sample, and there aren't going to be any white-collar workers or (probably) any wealthy people.

The *average* income might be comparable to the US *average* (not sure), but the distribution is going to be radically different.

Probably few over 60, though maybe not *as* few as one might expect –

But the big question of representativeness here is going to be: is there really a "viral dose" effect on severity of disease? (People keep suggesting this as a potential reason why healthcare workers seem to have had more bad outcomes than other people of comparable age.)

Also, infections vs. cases – did all these plants test everyone, even those without symptoms?

It's probably no worse than the Santa Clara thing, though, even with all that.

Twain on May 9, 2020 10:59 PM at 10:59 pm said:

Joshua,

Thank you for sharing! I'll take a look.

—

confused,

> But the big question of representativeness here is going to be: is there really a "viral dose" effect on severity of disease? (People keep suggesting this as a potential reason why healthcare workers seem to have had more bad outcomes than other people of comparable age.)

Here is a starting point for learning about this argument:

In terms of immunity, the lower the initial load of pathogen (virus, bacteria, etc.) entering the body, the higher the probability of the innate immune system quickly eliminating it. Therefore, the less SARS-CoV-2 that enters your body during a given period, the more likely it is your immune system can eliminate it before any infection or symptoms develop.

Twain on May 9, 2020 11:02 PM at 11:02 pm said:

Borked the link: https://en.wikipedia.org/wiki/Viral_load

Note: It is not the best source; immunology textbooks are a better start, like Janeway's Immunology, but they usually are not free (ugh!).

Joshua on May 9, 2020 11:42 PM at 11:42 pm said:

confused –

> The *average* income might be comparable to the US *average* (not sure), but the distribution is going to be radically different.

Yeah – that's what I was going for with representativeness – but sure, that was pretty sloppy.

confused on May 11, 2020 5:13 AM at 5:13 am said:

Twain: yeah, I understand the concept, I was just saying that (I think?) it's uncertain how significant the effects are in the case of COVID-19.

If there is a strong viral dose effect, we would expect to see a higher IFR among people infected in situations like meatpacking plants, hospitals, nursing homes, where they are exposed to a lot of virus for a long time, than among people infected by relatively brief contact while grocery shopping or whatever.

Anon on May 27, 2020 7:28 PM at 7:28 pm said:

I agree wholeheartedly. Sigificant confounding and bias going on here.

Reply ↓

Joshua on May 8, 2020 10:21 AM at 10:21 am said:

As we see with climate change, much of this boils down to people with an agenda, either deliberately or through ignorance, ignoring the full range of confidence intervals and confusing "projections" with "predictions" so as to confirm ideological biases. Much of the public simply doesn't look at modeling from the "all are wrong and some are useful" framework. And that is easy for advocates to exploit.

The sort of people who comment at this site are extreme outliers, and need to take that into account when they evaluate the interplay between complex modeling and the public discussion of highly polarized topics such as COVID 19.

Ferguson's competence, or lack thereof, should be evaluated within a larger context, where the skill of his modeling and the public discussion about his modeling are only overlapping (but not congruent) domains. The impact of his modeling is not simply a function of the skill of his models.

Reply ↓

Twain on May 8, 2020 10:30 AM at 10:30 am said:

Joshua,

> As we see with climate change, much of this boils down to people with an agenda, either deliberately or through ignorance, ignoring the full range of confidence intervals and confusing "projections" with "predictions" so as to confirm ideological biases. Much of the public simply doesn't look at modeling from the "all are wrong and some are useful" framework. And that is easy for advocates to exploit.

Well said, and I agree.

My only comment: Is it not the responsibility of the expert to prevent others from sensationalizing their claims? If, for example, I published a paper and someone took the results out-of-context or sensationalized them, I'd do my best to correct them. But perhaps in the era of the Internet, Mass Media, and Social Media, this just is not reasonable to expect anymore. I'm not sure myself.

Reply ↓

Daniel Lakeland on May 8, 2020 11:35 AM at 11:35 am said:

>Is it not the responsibility of the expert to prevent others from sensationalizing their claims?

Hell no, media publishes whatever it wants, and they defend that crap to the death. There is *no* agency on the part of researchers to cause any change, other than to sue one media outlet after another for defamation or some such thing. That'd immediately cost a couple billion dollars given the number of outlets. No way.

Reply ↓

Martha (Smith) on May 8, 2020 4:30 PM at 4:30 pm said:

I agree that it is not the responsibility of the expert to *prevent* others from sensationalizing their claims — that would be an impossible job. But I think it is the responsibility of the expert (together with the organizations, such as universities and government agencies, where the expert works) to try to *refute* such sensationalizing when it does occur. The organizations typically have some kind of PR unit, so this should not be an onerous task for them.

Reply ↓

Daniel Lakeland on May 8, 2020 5:13 PM at 5:13 pm said:

Agreed to a point. The PR unit doesn't know how to formulate responses, they can just push them out and defend them. If there is a "Gish Gallop" of bullshit, then it can take full time fighting it by the original scientists. This actually plays into the purpose of the Gish Gallop. The point is epistemic con-men win even if you engage them tirelessly fighting all their bullshit in a successful way.

The only real way to fight this kind of thing is political not scientific. The point is it's a propaganda infowar, and you have to just call it out over and over "this is a propaganda infowar" pointing to a mild body of evidence regarding disputed claims, and then hold the line once sufficient evidence is there to show it's all bullshit.

Reply ↓

Ellie Kesselman on May 8, 2020 10:40 PM at 10:40 pm said:

Martha,
That is a valid point. I am reminded of Brian Wansink. Cornell University's PR group published many press releases about the status of the situation, as both the media and other academic researchers investigated and released findings. At a minimum, the public relations groups at the institution with which the expert is associated can refute or clarify/qualify misleading claims that the media or others may be promulgating.

Reply ↓

Joshua on May 8, 2020 10:43 AM at 10:43 am said:

Twain –

> My only comment: Is it not the responsibility of the expert to prevent others from sensationalizing their claims?

Well, prevent might be a bit strong. But yes, the expert should lead with the uncertainties. They should lead with the caveat that a given projection is conditional, based on highly predictive and highly uncertain parameters.

There may be a problem where they are publishing for a professional audience with an implicit expectation that the caveats are understood, but that targeted communication makes its way to the broader public where the caveats are not understood. So we should have reasonable expectations that a public misunderstanding may arise from that misalignment. But yeah, experts should be better at getting out in front of how their work might be misinterpreted.

> f, for example, I published a paper and someone took the results out-of-context or sensationalized them, I'd do my best to correct them. But perhaps in the era of the Internet, Mass Media, and Social Media, this just is not reasonable to expect anymore. I'm not sure myself.

I think there is a reasonable level of expectation, and that what I consider reasonable in that regard is, often, not met. On the other hand, I often see what I consider to be an unreasonable level of expectation whereby scientists are blamed for people either willfully or deliberately misconstruing their work.

This is a big issue in climate science, where people reverse engineer from the phenomenon of "skepticism" to blame the existence of that phenomenon on the poor communication strategies of climate scientists. But such a viewpoint ignores what we know about basic mechanics of motivated reasoning – such that people who are ideologically so predisposed will use whatever climate scientists say, no matter how they say it, to confirm their underlying beliefs which are congruent with their political orientation.

Typically, these issues don't arise – and they wouldn't arise with COVID-19 modeling – except that the modeling necessarily overlaps with a highly politicized context of policy development.

Reply ↓

Twain on May 8, 2020 10:56 AM at 10:56 am said:

Joshua,

Again, all good points!

Reply ↓

SDE on May 8, 2020 10:51 AM at 10:51 am said:

A quick review of Ferguson's code settles the value of his forecasts completely. It's terrifyingly bad. No result should ever be trusted that comes from it. It can't be fixed, debugged, trusted in any way. Aside from some deeply concerning issues about reproducibility in the logs, and a complete lack of appropriate testing, the logical structure is so poor that the use of the word "structure" is not appropriate. As a former academic, I get that "research code" isn't always great, but not even research on the lowest stakes questions should be published on code like this. Also, this code stopped being research code when it started being used for advising policy. I literally cannot be too hard on it. I would immediately "contact HR to implement a development plan" to use a euphemism, if even a junior scientist or engineer on my team submitted code like this.

Reply ↓

Bob on May 8, 2020 11:53 AM at 11:53 am said:

Exactly, have your crappy code and write your crappy research papers. But once it's informing policy, it's not good enough. As for informing policy responses in a crisis like this. It's criminal.

I don't buy people crapping on about incentives. This guy was very well funded, and working on something of enormous importance. It doesn't occur to anyome to hire people who will code the damn thing properly?

I can't believe the excuses people are making for him, (hey how did that replication crisis happen, it's a mystery to me!)

You cannot seperate his model from the implementation of it. The model risk is enormous, the consequences are enormous.

Oh academic incentives! boo hoo!

Reply ↓

dhogaza on May 8, 2020 12:47 PM at 12:47 pm said:

I imagine that you're upset because his model projected that a lockdown directive might limit deaths to 20K, and the reality is that there are already 35K? You're upset because his poor code and model low-balled estimated deaths?

Reply ↓

Bob on May 8, 2020 1:04 PM at 1:04 pm said:

You can imagine all you like, but you're wrong. That makes you look foolish.

It's also very revealing of your own thought processes.

It's interesting to me that on a blog dedicated in large part to researcher bias, sees those very same biases play out in real time.

Reply ↓

Ellie Kesselman on May 8, 2020 10:32 PM at 10:32 pm said:

dhogaza
I'm kind of shocked to read your comment. Both SDE and Bob give well-supported reasons as to why Ferguson's model should never have been used as the basis for policy making.

Aren't we all saddened by lower death estimates than the reality turns out to be? That is a human emotional response, whether Ferguson and his academic institution's leadership had been better or not.

Reply ↓

dhogaza on May 8, 2020 10:42 PM at 10:42 pm said:

My comment to Bob was sarcastic.

Bob's comments regarding the C++ code are just stupid. The problem regarding seeding isn't as he described it, and his description of the comment thread makes clear he didn't understand the correspondence. I think it's clear he's just parroting what this "lockdownsceptic" person wrote, as the mistakes are essentialy equivalent.

SDE writes things that sound authoritative, but read closely. He's arguing from a position of personal authority. He's hand-waving. He doesn't actually give examples of code that demonstrates why he thinks it is so horrible that it is useless.

He also said at one point that the code is so bad it can't be discussed, verified, or debugged yet … the issue thread the Bob harps on shows that once the bug was made clear, a developer had debugged, fixed, and committed the fix within two days.

It's all politically-driven bullshit.

Reply ↓

dhogaza on May 8, 2020 10:46 PM at 10:46 pm said:

"Aren't we all saddened by lower death estimates than the reality turns out to be? "

To be blunt, no. Because I see no evidence that they care at all about the model projections for the lockdown scenario.

The attempts to discredit and trash the model stem from the fact that the lockdown was imposed in the first place. Why, if they'd understood the model was trash, there would've been no lockdown, and everything would've been just fine! That's the underlying storyline behind this.

Reply ↓

Vidur Kapur on May 8, 2020 10:59 AM at 10:59 am said:

On BSE, he predicted 50 to 50,000 deaths *by 2080*. His lower bound was 40 deaths and his upper bound was ~7,000 deaths for cumulative deaths by 2020.

He emphasised that the worst-case scenarios were far less likely than other scenarios and the paper suggests that the best estimate is 100 deaths to 1,000 deaths. See: https://www.nature.com/articles/nature709/ and https://www.nature.com/news/2002/020107/full/news020107-7.html

On bird flu, he was predicting the possible death toll in a scenario in which bird flu became highly transmissible between humans and retained its CFR of ~60%, which is something that infectious disease experts continue to worry about (see Amesh Adalja on Sam Harris' podcast, for instance).

His CFR estimate of 0.4% for swine flu was indeed used to calculate a worst-case scenario for swine flu by the U.K. Government. But it was clear that this was the CFR, not the IFR, and ultimately when calculating the CFR all you can do is look at global published data and update as more data come in.

Reply ↓

Andrew on May 8, 2020 11:04 AM at 11:04 am said:

Vidur:

Thanks for looking this up. That's helpful.

Reply ↓

Vidur Kapur on May 8, 2020 11:04 AM at 11:04 am said:

I'd also like to emphasise that Ferguson predicted 7,000 to 20,000 deaths in the UK if we went into lockdown. We're currently past 30K officially and excess mortality data suggest we may have had 50,000 deaths.

https://www.newscientist.com/article/2238578-uk-has-enough-intensive-care-units-for-coronavirus-expert-predicts/

https://www.reuters.com/article/us-health-coronavirus-britain-ferguson/uk-coronavirus-deaths-could-reach-7000-to-20000-ferguson-idUSKBN21N0BN

Reply ↓

Carlos Ungil on May 8, 2020 12:01 PM at 12:01 pm said:

So he got it wrong either way! QED. /s

Reply ↓

Funko on May 8, 2020 4:01 PM at 4:01 pm said:

:D

Reply ↓

Stephen M St Onge on May 11, 2020 4:12 PM at 4:12 pm said:

So let's see. Ferguson was low on deaths in lockdown by a factor of 1.5 to 2.5, so far. And it seems intuitively plausible that total deaths could ultimately double. So this means we should do what? Multiply his 500,000 worst case coronavirus by three to five, and get 1.5 million to 3 million deaths worst case? If so, then lockdown in Britain prevented around 96% of deaths, and saved the lives of about 2,900,000 Britons. I guess he should be a national hero: 'He prevented a disaster on the scale of the Great War.'

Or maybe we should say 'Since he blew the lockdown upper bound, we should conclude that lockdown has had no effect on the death toll, and excoriate him for recommending a bad policy!'? I admit I don't quite see the logic on that one, but it seems to be what is being suggested.

I have another suggestion. Instead of trying to prove one's superiority by sneering at someone, we might try to improve the modelling for next time.

"NAH!", as Steve Martin would say. Our inner monkey knows that nothing fundamental ever changes, so let's go with humiliating Ferguson, which automatically raises our status in the monkey troop.

Feh!

Reply ↓

Daniel Lakeland on May 11, 2020 4:19 PM at 4:19 pm said:

Just to be clear, I'm 100% sure Carlos' original was tongue in cheek. But St Onge's perspective aligns with mine, there's no question Ferguson saved a lot of people by convincing the UK to do something rather than their original plan of "do nothing much and let herd immunity occur"

I'm hoping the UK starts towards testing and tracing, and/or scheduled on/off cycles to break chains, but I confess to having no idea what they're up to over there.

Reply ↓

dhogaza on May 8, 2020 1:02 PM at 1:02 pm said:

Yeah, but the self-appointed expert code reviewers (note they don't criticize the actual model that the code implements, and almost certainly aren't qualified to do so) are ignoring that inconvenient fact.

They're focused on the fact that 500K haven't died, and apparently won't die.

Bob makes it clear:

"That picture has become less certain since the actual out-turn of deaths, the redundancy of the surge capacity that was built, and events in Sweden."

He obviously isn't referring to the model under-projecting deaths in the lockdown case … he's referring to the fact that under lockdown the no-lockdown projection hasn't been met.

And, as one could predict … "Sweden". Ignoring that deaths per capita in Sweden are 3.4 times higher than in neighboring, demographically, economically similar Denmark. About 10-14 days ago the death rate in Sweden was "only" about 2.8 times as high as Denmark. Ignoring that Sweden, concentrating efforts on isolating vulnerable people, hasn't been able to keep it out of nursing homes. Etc.

Reply ↓

Bob on May 8, 2020 1:07 PM at 1:07 pm said:

Nope, no he's not.

You're a walking example of why there's a replication crisis in the first place.

Reply ↓

Stephen M St Onge on May 11, 2020 11:50 PM at 11:50 pm said:

Concerning Sweden, Denmark, and all other countries, we need to know how deaths are being counted, and how accurately they are being counted, regardless of the criteria used for labeling something as a death from CCP Virus.

It was while trying and failing to get information on this that I found that the only accurate information I could get on Sweden is that the Somali immigrants have a much higher infection and death rate. than ethnic Swedes.

All this makes comparisons between countries very difficult.

Reply ↓

Moira Inness on September 27, 2020 2:55 PM at 2:55 pm said:

Never trusted scientists since they produced'the thalidomide drug which the claimed was safe to me prescribed for sickness in pregnancy. The thalidomide drug killed and maimed thousands of new born babies.

If Ferguson was so convinced his predictions were true why did he risk meeting up with his lover?

No justification for the lockdown, the pandemic is in fact flu epidemics that we experience every winter. It is criminal the actions he has initiated with his badly researched, (if indeed he did do any) and grossly extremely exaggerated predictions.

Reply ↓

Dale Lehman on May 8, 2020 11:17 AM at 11:17 am said:

Thanks for another reference point. I have no idea whether Ferguson is vindicated or not and I haven't looked at his COVID code so I am not defending it or his analyses. But at this point I am struck by how much the Fund article that Andrew has quoted is getting re-reported around the world. Bad news (pun intended) does indeed travel fast. It's time we recognize just how toxic the environment has become (another pun, this one unintended). We can belabor what responsibility researchers and journalists should have, but the truth is that news is anything that gets people's attention. And the work required to walk back something that is in error is far outweighed by the ease with which incorrect or misrepresentations can spread. Perhaps we should estimate the R0 for attention-grabbing headlines – I'll suggest is is >10 (no data, no code provided).

Reply ↓

Vidur Kapur on May 8, 2020 11:52 AM at 11:52 am said:

Andrew and Dale: no problem. I'd actually looked all of this up a few weeks ago because these claims have been circulating in the British media for a while now!

On COVID, Professor Ferguson estimated that we'd be unlikely to see >20,000 deaths under lockdown conditions. His range was 7,000 deaths to 20,000 deaths. We've had >30,000 officially confirmed deaths and Chris Giles of the Financial Times uses excess mortality data to estimate that we've had ~50,000 deaths. If anything, Prof Ferguson seems to have been optimistic so far, though he was spot-on when he predicted that we wouldn't see ICU overload under lockdown conditions.

And if you look at the March 16 report, his model predicted that under lockdown conditions, ICU bed use would peak in mid-April. That's exactly what happened.

https://www.newscientist.com/article/2238578-uk-has-enough-intensive-care-units-for-coronavirus-expert-predicts/

https://www.reuters.com/article/us-health-coronavirus-britain-ferguson/uk-coronavirus-deaths-could-reach-7000-to-20000-ferguson-idUSKBN21N0BN

Reply ↓

yyw on May 8, 2020 1:45 PM at 1:45 pm said:

Not exactly clear what assumptions were behind these numbers. His paper predicted at least 200k+ deaths if lock down lasts until end of June. The numbers you quoted above seem consistent with projections based on 5-month lock down in his paper.

Reply ↓

Vidur Kapur on May 8, 2020 4:23 PM at 4:23 pm said:

Thanks for your reply. The 250,000 death projection was for the "mitigation" scenario, which only involved isolation of cases, quarantine of the rest of the household, and "shielding" of the over-70s.

If you go to Figure 3B, which looks at the suppression strategy (which we're now following), the green line shows ICU use peaking in mid-April.

And if you look at Table 4, which looks at the idea of "intermittent lockdowns", the right column of the "total deaths" box shows all of the projected death tolls (for various "lockdown triggers") in the tens of thousands, the highest being 48,000 for an "on" trigger of new weekly ICU cases of 400 assuming an R0 of 2.6).

Reply ↓

jim on May 8, 2020 11:15 AM at 11:15 am said:

Thanks to everyone who provided context on this.

My conclusion: Exonerated of all charges, except maybe crappy code.

Reply ↓

Jonathan (another one) on May 8, 2020 11:19 AM at 11:19 am said:

And other than that, Mrs. Lincoln, how did you like the show?

Reply ↓

jim on May 8, 2020 11:36 AM at 11:36 am said:

I don't get your comment.

From a scientific standpoint, the evidence presented by commenters indicates that Ferguson has provided appropriate and responsible constraints for the forecasts in question.

How he's behaved with respect to public policy I'm not sure, and how politicians or the media have distorted his results I cant say, but I don't think that's the question in this post.

Reply ↓

Daniel Lakeland on May 8, 2020 11:46 AM at 11:46 am said:

I agree.

Reply ↓

Bob on May 8, 2020 11:59 AM at 11:59 am said:

The point is the model IS the code.

His predictions are non deterministic given a random seed. There is variation in his predictions that

1. He cannot explain, and

2. He had absolutely no interest in tracking down and correcting. The results should just be averaged.

But while not fixing the issues, they were adding more 'featues' to the model.

Reply ↓

dhogaza on May 8, 2020 1:08 PM at 1:08 pm said:

"His predictions are non deterministic given a random seed"

Did you really mean to say a random seed?

Anyway, I do share your concern. It is really unfortunate that his 20K projection for deaths under lockdown was far too optimistic, given that the UK now has 35K dead and it's not over yet. Thanks for pointing out that his projections were too low.

Reply ↓

Bob on May 8, 2020 1:28 PM at 1:28 pm said:

Try reading what's on github, you might stop making yourself look foolish.

Reply ↓

Zhou Fang on May 8, 2020 2:43 PM at 2:43 pm said:

Github says they fixed the issue, which alluded to when the RNG was being reinitialised, and it doesn't affect the validity of the model?

https://github.com/mrc-ide/covid-sim/commit/616a814ee94497c540c9d7b94895a8f17584eb35

So, it didn't matter?

dhogaza on May 8, 2020 4:54 PM at 4:54 pm said:

Zhou

"So, it didn't matter?"

An earlier version of the model allowed one to hardwire a seed value, making the model deterministic. Useful for validation and debugging.

They changed how they were seeding the rand function and introduced a bug which made it impossible to hardwire a seed value and have it used throughout the model run, removing determinism. It's not clear how long the bug was there. I suppose one could look at the change log and see if one really felt it was crucial. My personal guess is that this change came long after the code was originally written 13 years ago.

It has been claimed that this "post-cleanup" code up on github isn't the code used to run the projections.

So, THIS BUG MAY NOT HAVE BEEN PRESENT IN THE VERSION USED TO RUN THE PROJECTIONS.

Gosh, that doesn't really fit their narrative, does it?

No, this does not impact the validity of the implementation of the underlying model (which is what folks should focus on).

Elsewhere it has been said …

" It's so bad that you can't even have a conversation about it, verify it, debug it …"

Someone debugged it. Even fixed it, and committed the change, with a rather detailed commit message. Apparently the above statement isn't strictly true.

More generally, poorly written code is more difficult to work on. Harder to debug. Harder to verify. Not impossible.

In some ways easier to talk about, though, I don't know what that person was going on about, look at all the discussion here!

somebody on May 8, 2020 3:45 PM at 3:45 pm said:

This is just a lie. You're referencing a blog post, which linked to an issue tracker, which brought up an issue with a new feature for reading in pre-built network files re-initializing the random seed in between runs. The issue was documented on a repo which was not the code used to generate the predictions, but rather an extension of the modeling concept, was introduced in a new feature exclusive to the extended code, and has actually been both explained and fixed. This has absolutely nothing to do with his predictions. Did you actually read the issue in question? Your number 2 is, as far as I can tell, just a lie, just something that YOU MADE UP.

Reply ↓

dhogaza on May 8, 2020 5:01 PM at 5:01 pm said:

somebody:

"The issue was documented on a repo which was not the code used to generate the predictions, but rather an extension of the modeling concept, was introduced in a new feature exclusive to the extended code, and has actually been both explained and fixed. This has absolutely nothing to do with his predictions."

Ha! Wish I'd read that before I posted to the subthread above!

I said:

"It has been claimed that this "post-cleanup" code up on github isn't the code used to run the projections.

So, THIS BUG MAY NOT HAVE BEEN PRESENT IN THE VERSION USED TO RUN THE PROJECTIONS."

So I was right, huh? Why am I not surprised. Thanks for digging into it so you could describe precisely what was up with that.

As far as the lack of maintainability of the supposedly unreadable code base, looks like people are extending it, no problem. And, oh gosh, a bug was introduced while writing a new feature. And then debugged. And fixed. And committed with an informative commit message.

"Did you actually read the issue in question?"

I think you understand that they're not interested in understanding the code, simply slinging arrows at it.

Reply ↓

Jonathan (another one) on May 8, 2020 12:14 PM at 12:14 pm said:

What Bob said.
Sure, he's not guilty of forecasting outside of really really wide bounds. But if those forecasts are based on models which can't be vetted, why should anyone care? "I used my random number generator, plus some other code, 10,000 times and, assuming my random number generator, plus some other code, is the true model, I get a 95% CI of between 50,000 and 98,000." Why is only the last clause of the statement worth vetting?

Reply ↓

Joshua on May 8, 2020 12:21 PM at 12:21 pm said:

Jonathan (another one –

> But if those forecasts are based on models which can't be vetted, why should anyone care?

A question for you: Do you care that people are predicting widespread economic and other social harms as a result of government mandated social distancing? If so, why do you care?

Reply ↓

Jonathan (another one) on May 9, 2020 1:31 PM at 1:31 pm said:

Not sure I understand your question.

Do I care what some people are predicting? I guess that depends on whether people are listening to them and taking actions based on what they're hearing. Otherwise, definitely not.

Do I care whether their predictions have a good underlying basis? Not unless the previous test is satisfied. But if it is, it's really all I care about.

Reply ↓

Joshua on May 9, 2020 6:43 PM at 6:43 pm said:

Sorry for being obtusd.

I have observed a lot of climate "skeptics" claim that climate models are worthless because they can't be "validated" and because they project in a range which then gets twisted into wrong "predictions" if the highest end of the ranges aren't reached.

But then those same "skeptics" rely on unvalidated modeling all the time to confidently predict economic disaster if emissions are mitigated.

I see something similar in the COVID-19 wars. Rightwingers say that projections from models are just scare-mongering because they could be wrong, but then they turn around and say that government mandated social distancing has caused economic harm – even though they have know way of actually disaggregating the impact of the mandates from economic impacts from massive death and illness.

My basic point is that people tend to have unrealistic standards for models when they don't like what the models project.

IMO, the point is that models can help to clarify uncertainty – and as tools that are always wrong but sometimes useful, they can help people to make decisions in the face of extreme uncertainty. I'm leery of people dismissing the output of complex models by having unreachable standards. Not saying that applies to you – just doing my due diligence.

dhogaza on May 8, 2020 1:11 PM at 1:11 pm said:

"I used my random number generator, plus some other code, 10,000 times"

Monte Carlo simulations are used everywhere.

Reply ↓

Jonathan (another one) on May 9, 2020 1:34 PM at 1:34 pm said:

No one disputes this that I'm aware of, not even 'Sue Denim.' The critical clause in my statement was 'assuming it is the true model.' That's the issue, not Monte Carlo analysis. And if the code is sloppy, it's very very difficult to figure out whether the model is 'true enough for government work' or not.

Reply ↓

dhogaza on May 9, 2020 1:51 PM at 1:51 pm said:

Well, I've dug through a bit of the code and didn't find it sloppy or hard to read at all, and in about 15 minutes managed to glean some high-level information about it.

See my post from earlier this morning.

Maybe I'm just better at reading code than some of our self-proclaimed experts who've shown up here.

"Sue Denim" showed that "she" couldn't even understand what was being said in the github issues thread documenting the seeding issue. I checked it out. It was perfectly obvious. It was also obvious that the model was deterministic in SP mode if run the way the developers do for testing. Bob swallowed "Sue Denim"'s claims uncritically, which could be because he was lazy or could be because he doesn't have the skills to validate "her" bogus claim.

SDE claims the model can't be discussed because it's written so badly, but again, I spent about 15 minutes and gleaned quite a bit about how the basic agent-based mechanism works inside the model. Code was clean, variables and other entities assigned informative names, etc.

So it's all bullshit on their part.

But regardless, your comment is meaningless. There's no such thing as "the true model". All models are approximations. Some are better than others.

jim on May 8, 2020 11:34 AM at 11:34 am said:

So what's the big hoo ha about releasing code?

People seem to want to accept/reject model forecasts on the basis of model composition, rather than bother with the pesky little step of actually finding out of the model is accurate.

No model should be used for public policy until it's results have been tested against reality and verified to be accurate within a useful error range.

Reply ↓

Daniel Lakeland on May 8, 2020 11:41 AM at 11:41 am said:

You can't test models of extreme events against reality. Suppose I have a model of tsunamis which says that 15m tsunamis occur on average every 600 years in the vicinity of Fukushima, am I to wait around for 2400 years so I can see that about 4 of those occur before I'm allowed to weigh in on whether we should build a nuclear reactor at Fukushima?

The only thing you can do is look at the proposed mechanism and see if it makes sense. You can maybe use the same model to predict some other events, like for example building a physical scale model in a lab, and trying to use the code to predict what happens there… But how many scale models of disease transmission in the modern world, with airplanes and mass gatherings for worship and soforth are you likely to have?

Reply ↓

> Martha (Smith) on May 8, 2020 4:43 PM at 4:43 pm said:
>
> Good point.
>
> Reply ↓

Bob on May 8, 2020 12:03 PM at 12:03 pm said:

Nope, they want models that are implemented correctly, so that the predictions are a result of your model and not random errors introduced because you can't be bothered to code your model correctly.

Reply ↓

Zhou Fang on May 8, 2020 2:20 PM at 2:20 pm said:

Then build your own model and demonstrate it gives different results?

Reply ↓

Martha (Smith) on May 8, 2020 4:45 PM at 4:45 pm said:

Both Bob and Zhou Fang make good points.

Reply ↓

John Williams on May 8, 2020 12:24 PM at 12:24 pm said:

"No model should be used for public policy until it's results have been tested against reality and verified to be accurate within a useful error range."

What about the conceptual models that ordinary people (including politicians)use to think about things? At least coding a model makes you be explicit, and, at least in my experience, often shows that your conceptual model can't be right.

Reply ↓

SDE on May 8, 2020 12:40 PM at 12:40 pm said:

Look, the code that generates the forecasts is so bad that it is totally useless. This isn't nitpicking, or cosmetics. It's fatally badly written, can't really even be shown to be wrong. It's so bad that you can't even have a conversation about it, verify it, debug it, or even rationally relate input to output. Since it is open sources, it's a matter of public record. Don't trust anything from it. Right or wrong forecasts don't enter into it. Fortune tellers are right sometimes too.

Reply ↓

Bob on May 8, 2020 1:35 PM at 1:35 pm said:

You're wasting your time.

They've decided that anyone who points it out has 'an agenda', and they're now fighting the good fight and won't concede you have a point. I mean, you just have to read the issue tracker but that's too much work.

Reply ↓

SDE on May 8, 2020 1:55 PM at 1:55 pm said:

Or glance at the code and know more about code than CS 101 . . .

Reply ↓

Bob on May 8, 2020 2:20 PM at 2:20 pm said:

If anyone ever asks again what caused the replication crisis I'm goint to point them to this blog post.

Reply ↓

dhogaza on May 8, 2020 9:57 PM at 9:57 pm said:

Bob

"If anyone ever asks again what caused the replication crisis I'm goint to point them to this blog post."

You really want them to see that you don't actually understand what the replication crises refers to?

GCB on May 8, 2020 2:50 PM at 2:50 pm said:

Come on, guys. Reciting points taken from a clearly agenda-driven blog indicate that you have an agenda. The code in question is not the code used for the policymaking. It's being written by Microsoft and Github engineers (and others) who state their code is under active development and make pretty clear that it should be used by people unfamiliar at their own peril. Enough of the "useless code" and "code is the model" canards.

Reply ↓

SDE on May 8, 2020 5:44 PM at 5:44 pm said:

To make (another) attempt to ground this in fact, Github, as is standard for git, has the virtue of showing you commit history. The code is mostly not new, and Ferguson is all through it. I'm not reciting from a blog, but have rather looked at the code from the perspective of someone who has written lots of production code at top tech firms, and also used to be an academic. Inferring an agenda based on criticizing the code, a matter of public record, is simply not sound. Like most of the serious developers on this thread, I've kept my criticisms strictly to the technology, which I understand quite well, thank you.

I've hired from Microsoft. Not all of their candidates are great, but I've never had a candidate who showed me code like this. They obviously aren't far into any refactor, and I expect they'll generate something good. Many engineers at my major tech firm place of employment, engaged this on our tech chat channel including some very, very senior engineers. This code caused scandal, and not one engineer, regardless of a wide range of viewpoints (agendas if you don't like them), had a kind thing to say about the code. The kindest thing said was "terrible, but not actually the worst I've seen".

Reply ↓

Zhou Fang on May 8, 2020 6:46 PM at 6:46 pm said:

Liar.

Zhou Fang on May 8, 2020 6:46 PM at 6:46 pm said:

The Github commit history starts 17 days ago.

GCB on May 8, 2020 8:33 PM at 8:33 pm said:

I like you. You're funny!

GCB on May 8, 2020 8:49 PM at 8:49 pm said:

Should clarify: SDE is comical. Zhou Fang is just more blunt than I would have been.

dhogaza on May 8, 2020 10:08 PM at 10:08 pm said:

"Inferring an agenda based on criticizing the code, a matter of public record, is simply not sound."

OK, point to a file and tell us exactly what you find so egregiously bad about it the code.

Xoogler on May 9, 2020 5:23 AM at 5:23 am said:

I find it hilarious that you're trying to debate via intimidation/argument from authority via making up lies about programming, despite this blog being followed by a fair number of technically competent people who are at least capable of reading Git.

Since you're so excited about arguing from authority, which "major tech firm" are you from, exactly, Mr."SDE?" Google doesn't have "SDE" as a title, also there's no "tech chat channel." SDE isn't a title at Netflix either. I highly doubt you're from Microsoft.

What you said doesn't sound like Facebook's culture or technical standards either, though it's just barely on the edge of possibility. So maybe you work at Apple, which is too secretive for me to know much about their internal processes?

In case you don't want to disclose, right now my odds are 5% Apple, 10% one of the other major firms, 25% 4th-tier-firm-that-thinks-it's-"major", 60% entirely made this story up.

Hope this is helpful for non-tech people trying to navigate his (likely) BS arguments from authority.

dhogaza on May 9, 2020 11:28 AM at 11:28 am said:

SDE earlier said:

"Look, the code that generates the forecasts is so bad that it is totally useless. This isn't nitpicking, or cosmetics. It's fatally badly written, can't really even be shown to be wrong. It's so bad that you can't even have a conversation about it, verify it, debug it, or even rationally relate input to output"

We already know it can be debugged, as someone quite quickly debugged and fixed the seeding error that caused different results results in SP mode if you first generated a network, saved it, restarted with the saved network vs. just run it in one step. If it was run in one step the result was deterministic, and apparently this is how the development team normally did their testing, so had not run across the issue. Nothing burger.

So I decided to spend a few minutes poking around to see if it's so bad you "can't even have a conversation about it, verify it".

The code to change the state of individuals in the model can be found starting in the beginning of this file:

https://github.com/mrc-ide/covid-sim/blob/master/src/Update.cpp

Looking at this it's immediately clear that this indeed is an agent-based model, working at the granularity of an individual. Three states, so it is an SIR model (though they use the terminology Susceptible, Latent, and Immune). The functions that move an individual from one state to another are straightforward enough, and from looking at the code that moves an individual you can easily see that it tracks the spread of the infection geometrically from an initial infection point as a circle, with the radius of that circle growing if the newly latent individual's household lies outside the currently infected zone.

It is claimed that the source code is so horrible that it is impossible to "even have a conversation about it". Clearly that's bullshit. I found out all of the above with a quick skim of the file. It's easy enough to read and follow the code. One of the most important factors that impacts the readability of the code has to do with the naming of variables, constants and functions and the naming used in the code is very easy to understand.

So I'm calling bullshit on the code quality claims.

Now I went about this ass-backwards because I was just interested in the claims regarding the quality of the code.

If I wanted to know how the model worked, I wouldn't start by looking at the code. Ferguson's a researcher, surely he and his team have published papers on epidemiological modeling that would help to understand the agent-based approach he takes. And when looking at the code, if learning it was really by goal, I'd start with the definitions of the data structures, state representations, and parameter definitions rather than code that manipulates them.

But I still managed to learn an interesting amount of high-level stuff just by skimming one file for a few minutes. I skimmed it for less time than I'm taking writing this post.

Now, do you need such a complex model to make high-level decisions?

That's a totally different discussion.

dhogaza on May 9, 2020 11:40 AM at 11:40 am said:

SDE:

And to rub salt into the wound, two weeks ago I had never even heard of an SIR or SEIR model, and I'm not a modeler in any field (though the concept of an agent-based model is trivial, of course). I admit to having spent some time looking into epidemiological modeling the past couple of weeks, motivated of course in a general sense by the times we are living in, and more specifically with James Annan's work using a Bayesian framework with an SEIR model (really a six box S E0 E1 I0 I1 R model, which he poached from an epidemiologist's online blog post) to predict the future course of the epidemic and realistic values for R0, and the change in Rt when the lockdown was put into place.

So I've been doing some background studying before plunging into the Ferguson model's code. I humbly suggest that if your goal is to understand that model, diving into the source code isn't the place to start.

But of course we know that your goal isn't to understand the model. It's to piss on it.

Kevin on May 8, 2020 2:17 PM at 2:17 pm said:

I'll take John Carmack's opinion over yours.

"That was my fear — what if the code turned out to be a horror show, making all the simulations questionable? I can't vouch for the actual algorithms, but the software engineering seems fine."

https://twitter.com/ID_AA_Carmack/status/1254787975459045381

Reply ↓

SDE on May 8, 2020 2:45 PM at 2:45 pm said:

Cool. Or better, look at the code yourself.

Reply ↓

dhogaza on May 8, 2020 10:11 PM at 10:11 pm said:

I've skimmed it. I was surprised not to see the kind of horror show I've been told I'd find.

Time for you to educate us by some examples from specific files that justify the extreme reaction we're seeing from you and others.

Reply ↓

dhogaza on May 9, 2020 11:31 AM at 11:31 am said:

I've looked at the code, and posted about my experience above. You're bullshitting. Just a few minutes reading one of the more basic files made that clear.

Reply ↓

Dave on May 8, 2020 11:58 AM at 11:58 am said:

Without commenting on the specifics of Ferguson, I think we may just have to accept that pandemics are quite difficult to predict. Even modest changes in parameter values can make an order of magnitude difference in conclusions, it seems. It's also a dynamic system in which people react voluntarily and through government actions to avoid infection. The more dangerous, the larger the behavioral change.

I think the quick work Cochrane did to include endogenous behavior that reacts to the perceived danger is at least a better mental model, even if I still wouldn't trust something like it to make precise predictions:

https://johnhcochrane.blogspot.com/2020/05/an-sir-model-with-behavior.html
https://johnhcochrane.blogspot.com/2020/05/dumb-reopening-might-just-work.html

Reply ↓

Terry on May 8, 2020 4:44 PM at 4:44 pm said:

+1

Reply ↓

Old Timer on May 8, 2020 1:33 PM at 1:33 pm said:

My old man used to day, "Son, there are three types of forecasters – @#%$ forecasters, accurate forecasters, and forecasters I can use… "
(CI, 95%)

Reply ↓

Dale Lehman on May 8, 2020 1:41 PM at 1:41 pm said:

Andrew: I humbly request that you retract this post. Debating Ferguson's model and role is certainly worthwhile and I am very interested in what people have to say about that. However, your post is about John Fund's article, and his failure to seriously address any of the concerns raised about his own work and credibility is no less serious than the issues you have repeatedly pointed out about researchers unwilling to admit error or seriously engage in discussion of their work. I am fed up with journalists (ok, maybe he is not a journalist but some other sort of media personality) believing they should get a free ride because they are investigating somebody else's work.

Of course, it is Fund that should retract his column. But, I'd like to see you retract or end this discussion under the guise that it has anything to do with his article. You've blogged about Ferguson's studies before – can't we continue that, but disassociate it from Fund's blatant editorializing?

Reply ↓

Andrew on May 8, 2020 2:02 PM at 2:02 pm said:

Dale:

I added a P.S. to clarify that explanation 1 seems to have some truth to it.

Reply ↓

Joshua on May 8, 2020 2:13 PM at 2:13 pm said:

Not that you should care about my opinion about what you should post about…but instead of retracting, why not just follow up with a post about how bad Fund's characterization of the Imperial College model was (and/or his history of mischaracterizations)?

Reply ↓

Andrew on May 8, 2020 2:23 PM at 2:23 pm said:

Josh:

There's room for something more statistical on the general challenge of reporting on uncertain forecasts in an environment where, on one hand, pundits and journalists don't want to miss the big story (FOMO!) and, on the other, there's pressure not to "panic" (recall our discussion from a couple months ago of he-who-shall-not-be-nudged).

On the other hand, I have some books to write, and here I am spending time in the comments section . . .

Reply ↓

Joshua L Brooks on May 8, 2020 2:31 PM at 2:31 pm said:

Andrew –

> There's room for something more statistical on the general challenge of reporting on uncertain forecasts in an environment where, on one hand, pundits and journalists don't want to miss the big story (FOMO!) and, on the other, there's pressure not to "panic" (recall our discussion from a couple months ago of he-who-shall-not-be-nudged).

Sure. But I come here from the climate blogosphere – and it's rather striking just how similar the parallels are w/r/t models, statistical analysis, public opinion, and tribalism. I also come here from the motivated reasoning blogosphere (Dan Kahan's defunct blog), and again the parallels are striking. There's a lot to dig into.

> On the other hand, I have some books to write, and here I am spending time in the comments section . . .

I appreciate your input at the level you have to spend time on it, and get that this stuff can be real time sink.

Reply ↓

John Fund on May 8, 2020 2:29 PM at 2:29 pm said:

Hmmm…so now I have a sudden "history of mischaracterizations." Thrown out on the table, just like that.

I await a comprehensive list. Sorry, Media Matters run by egregious David Brock is NOT a reliable source. Ask any fair-minded person in the media about that.

Reply ↓

Joshua on May 8, 2020 2:31 PM at 2:31 pm said:

John –

> I await a comprehensive listJust

Reply ↓

Joshua L Brooks on May 8, 2020 2:35 PM at 2:35 pm said:

Why not just begin with a few that have been described in this thread?

For example, why did you effectively truncate the uncertainty range of the Imperial College projections, and why did you not write about the fact that they were projections that were conditional, and based on incorporating into the confidence interval parameters that the authors clearly indicated were highly unlikely (such as no social distancing whatsoever)?

Reply ↓

Dalton on May 8, 2020 3:03 PM at 3:03 pm said:

I see you using "fair-minded", why not "fair and balanced"? I hear there are some "fair and balanced" people in the media.

I also think it's hilarious you're accusing others of ad hominem attacks followed up immediately by an ad hominem attack. Thanks for the chuckles!

Reply ↓

Linch Zhang on May 9, 2020 3:53 AM at 3:53 am said:

John, I can see how you are justifiably frustrated that this blog's statistically literate readers appear to be piling on you. However, your claim is that Ferguson is bad about making predictions because the actual numbers were nowhere near his upper bound in his 95% CI. This isn't just *bad* evidence for the accuracy of his forecasts, this is the opposite of evidence.

Structurally, here are the arguments you've presented at the beginning of your post:

"For i from 1 to 5:
(In X_i year, Ferguson predicted an upper bound of Y_i people will die. Less than Y_i people died}.

Thus Ferguson is WRONG!!11!!!
"

I see no reason to believe that this is even remotely valid as an argument. Like if his lower bound was too high, or if his upper bound was too low, okay then he was wrong. and overconfident Or maybe if he made 200+ predictions, and results always turned out to be lower than his 97.5th percentile, then you can say he's underconfident. But all I see so far is correctly calibrated predictions?

It's unfortunate but fine that you're wrong here. Plenty of journalists don't understand statistics, and I can see what you did as an honest mistake (maybe you missed reading the lower bound of the CI?). But since we're all reasonable adults here, I hope you can retract your article and issue a public apology to Nell Ferguson, as well as this blog's host and audience for wasting our time. :)

In turn, I will apologize on behalf of other critics of your character for casting aspersions on what is no doubt an honest mistake.

https://twitter.com/LinchZhang/status/1258765951257866244

Reply ↓

Terry on May 8, 2020 4:51 PM at 4:51 pm said:

John Fund = dishonest journalist.

Giving credence to anything he says is a mistake, I think. It might be ok to use his column as a jumping off point to look more carefully at a topic, but we should have strong priors against giving his writing significant weight.

Reply ↓

yyw on May 8, 2020 2:04 PM at 2:04 pm said:

I think a PS note on Fund's selective quoting would suffice. Some of Ferguson's predictions were clearly out of line. In the Guardian article, he gave an estimate of up to 200mil deaths for bird flu. There's no excuse for that (unless the reporter for that article completely misquoted him).

Reply ↓

Dalton on May 8, 2020 2:32 PM at 2:32 pm said:

Why is there no excuse for that prediction?

David Nabarro, Chief avian flu coordinator for the United Nations, stated it could have killed up to 150 million people worldwide. Note the "up to", the range was 5 to 150 million. The same figure the World Bank was using. All of this was predicated on the big "if" it mutates into a form that was readily transmissible among humans.

https://www.nytimes.com/2006/03/27/world/americas/the-response-to-bird-flu-too-much-or-not-enough.html

I don't see a problem with these predictions predicated on plausible hypothetical. As a society we should, to a certain extent, plan for somewhat unlikely but plausible disasters. I really don't expect my house to burn down and yet I still have insurance. And to pay for that insurance, I had to estimate what I could potentially lose if a catastrophe did hit.

Reply ↓

Charlie on May 8, 2020 3:35 PM at 3:35 pm said:

Seems like a lousy project if the guess is between n and like 7 orders of magnitude higher. At that point it just seems like a dart throwing chimp? That estimate doesn't seem to provide much help.

It seems like humans would see the far end (150 million) and panic, rather than thinking about the whole range of possible number of deaths.

Reply ↓

Terry on May 8, 2020 4:53 PM at 4:53 pm said:

+1

Reply ↓

Martha (Smith) on May 8, 2020 5:02 PM at 5:02 pm said:

"Seems like a lousy project if the guess is between n and like 7 orders of magnitude higher. At that point it just seems like a dart throwing chimp? That estimate doesn't seem to provide much help."

Yes — but that level of precision may be the best we can do. What we do need to do is acknowledge the inherent limitations in getting the precision we would like to have.

"It seems like humans would see the far end (150 million) and panic, rather than thinking about the whole range of possible number of deaths."

Yes — and we often fail (when communicating scientific topics) to take into account human nature, such as looking at the worst case scenario.

Putting both together: We (as a society) still fall short on effective communication of uncertainty, and on effective education on accepting and dealing with uncertainty.

Reply ↓

Zhou Fang on May 8, 2020 5:08 PM at 5:08 pm said:

I think the intended effect though is to spur the authorities to fund efforts to gather more data and in the meantime take steps, such as slaughtering chickens in the case of avian flu. You have to note that these projections were made with very few known cases.

Reply ↓

Phil on May 8, 2020 7:39 PM at 7:39 pm said:

"Yes — but that level of precision may be the best we can do. What we do need to do is acknowledge the inherent limitations in getting the precision we would like to have."

Exactly. If you're uncertain by 7 orders of magnitude then you should say so. What's the alternative? Pretend certainty you don't have?

If there's a reason to believe the upper bound is unreasonably high — no way in hell could this kill 200 million people worldwide — then sure, criticism is deserved. You can't just be inflating the upper limit for no reason. But if there were plausible conditions under which 200 million could die, someone can't be criticized for stating that. Or at least they shouldn't be.

Reply ↓

confused on May 8, 2020 7:54 PM at 7:54 pm said:

The 200 million number in itself is very shaky, because it's just derived by taking the estimated deaths from 1918 flu pandemic and adjusting for the increase of world population since then. There's no reason to think that the same % of deaths would apply to a different pandemic flu virus, and probably many people died in 1918 who would have survived with access to modern medicine. (Apparently there were many deaths from secondary bacterial pneumonia, which antibiotics could probably have greatly reduced; there were no ICUs in 1918; etc.)

But you're correct in general.

yyw on May 8, 2020 3:47 PM at 3:47 pm said:

How likely is the mutation? 1 in 100 chance, 1 in 1 million, 1 in a billion, or …? Anyway, the following is the Guardian article: https://www.theguardian.com/world/2005/sep/30/birdflu.-jamessturcke

You can judge for yourself if their estimates were way out of line.

Reply ↓

Zhou Fang on May 8, 2020 4:00 PM at 4:00 pm said:

At the time Ferguson made that statement there were about 100 confirmed cases of avian flu. Port that over to covid19 and you are basically asking Ferguson to make an estimate of the total global number of cases on Jan 20th. It seems ridiculous to expect that sort of precision at that point in time.

Reply ↓

Stephen M St Onge on May 12, 2020 12:50 AM at 12:50 am said:

The mistake is not understanding what is being said, not in the number.

Here's a statement: "Up to half a billion people could die in the next month." That's absolutely true. There are more than half billion people alive (somewhat over 7 billion currently), and they COULD all die in a month. Maybe the Sweet Meteor of Death finally impacts.

But the important part of my statement was not the number, it was the word "might." It's too vague to be falsifiable. The statement won't turn out to be wrong, because it isn't definite enough to be wrong.

Ferguson's statement about bird flu was that if it turned as deadly as the Spanish Flu, the world could experience the same percentage of the population dead as the world experienced in 1918-1920. Is that true? Beats me. In some ways we are worse off, with increased and faster travel making stopping the spread of an epidemic much harder. In other ways we are better off, with much improved communications.

IF a flu strain as infectious and deadly as the Spanish Flu comes along, and if it gets out of control, it seems at least plausible the same percentage of world population could die. We have 4.3 to 5.2 times as many people alive today, so pick your favorite SWAG as to flu deaths in 191-20 (17-100 million people), and multiply by both 4.3 and 5.2 times to get a range.

As I said in a comment above, what's wrong is taking the result of that exercise as a prediction. It's an analysis of how bad things could get, given certain assumptions. If you don't like the idea of seventy-three million to half-a-billion flu deaths, it suggests we take steps _now_ to prevent that outcome.

But people like predictions. So we will most likely continue to see them reported, and they will all be bullshit (except for the one I made in this sentence).

Learn to read VERY critically, that's my recommendation.

Reply ↓

John Fund on May 8, 2020 2:17 PM at 2:17 pm said:

I find it somewhat amusing that within 90 minutes of my arriving to this debate over my article, Dale Lehman is already asking that the entire post should be retracted. That's a mild way to justify a form of censorship along with an ad hominem suggestion that maybe I'm "not a journalist." Censorship and ad hominem are weak, weak responses in an admirable forum like this.

Please recall that too many of the web's gatekeepers were tempted to stamp out the free debate that helped alert Americans to the threat of the virus in the first place! Examples are abundant. Some want to require conformity with the judgment of expert institutions, even when many of those institutions themselves woefully misjudged the situation months or weeks ago.

Reply ↓

dhogaza on May 8, 2020 2:26 PM at 2:26 pm said:

I don't think it should be retracted.

Indeed, allow me to hand you a larger virtual shovel so you can continue digging an ever-deeper hole.

Reply ↓

Dale Lehman on May 8, 2020 2:35 PM at 2:35 pm said:

This will be my last comment about this. Retraction of a post from a blog is not censorship – particularly when your writing illustrates many of the practices that have been discussed on this blog, and which have been suitably criticized. Not for the substance of your views, but for sensationalizing, misrepresenting, and distorting facts. Had the post been about those things, then it would have been fitting for this blog. But the post had been presented as relevant to Ferguson's work. If the best relevance is Andrew's reference to his point #1, then it hardly needed your column to prompt that discussion.

You have plenty of outlets for your views (which I had no awareness of prior to this post). This blog has generally been free of trolls – it is the only place I am willing to use my real name. You apparently see any disagreement as censorship – it reminds me of an era when Communists were seen under every bed. It appears you thrive on riling people up. I guess getting attention is a sign that you are doing your job. I will not provide any more fodder for you.

Reply ↓

Nicolas Bray on May 11, 2020 4:14 PM at 4:14 pm said:

I guess blathering about censorship was easier than admitting that you were completely full of shit, huh.

Reply ↓

Simon on May 8, 2020 2:19 PM at 2:19 pm said:

+1

Reply ↓

Dalton on May 8, 2020 2:13 PM at 2:13 pm said:

Here is me way out on that limb, but I don't think Ferguson's original projections have been demonstrated to be wrong. The top line number for deaths in the United State he predicted was 2.2 million. We've seen various estimates of the IFR bandied about in previous threads, but (sorry Phil) I think an age-averaged IFR of 0.5 – 1% is not a bad guess. That would put it about 10 to 20 times as deadly a the flu. We've also seen R-naughts in the range of 2-4 and which from my lay reading suggest this thing won't slow down until something 2/3 of the population has been infected and either recovered or died. 328 million * 2/3 * 0.005 = 1.1 million dead and 328 million * 2/3 * 0.01 = 2.2 million dead.

(Aside, the other Imperial College model which was linked here and I think was somewhat well received has IFR estimates for each country in the range of 0.8 to 1%: https://mrc-ide.github.io/covid19estimates/)

That makes me think that Ferguson's original no intervention projection of the total number of deaths is plausible. I especially think it was plausible way back in mid-March when the only data we really had were cases beginning to sky-rocket in Italy and the example of China's extreme lockdown the only thing that slowed the virus down once it had reached a certain critical mass.

So then the question is one of timing and the effect of non-pharmaceutical interventions. Figure 1A in the original report does not show the peak in deaths occurring until mid-June in the UK and late-June in the USA. The peak in deaths in the UK and USA (after major interventions were undertaken) was something like 900 per day in the UK around April 14 and around 2,200 per day in the USA around April 20th (the USA really hasn't seen much of decline from that). That's around 0.7 deaths per 100,000 and around 1.4 deaths per 100,000 per day in the UK (populations of 328 and 66 million respectively). Which again is ballpark correct for that figure.

Then for the question of interventions. The most severe thing they looked at was "general social distancing" which I don't think means anything like China style lockdown. I don't even think they meant people not going to work. Under that scenario for the United States, which I think we largely adopted, they predict that critical care beds occupied would peak out at around 4 per 100,000 in the United States in early May (Figure A1 panel B). According to worldometers there are 16,000 or so cases classified as critical right now, which is right around 4 per 100,000.

So they wrote atrocious code. So what? So do I. I'm a statistician and biologist who knows a little something about coding. I try to make it efficient, I definitely make it replicable, but on the balance I know my code is terrible. I write my code to solve the scientific problem I'm trying to crack. (Seriously I've got a five-thousand line Stan model going right with a ton of copy-pasted variables that I refer to manually. I could probably find a more clever way to store these and refer to them, but all I know i that Stan won't allowed ragged arrays and I know where I want to get. It's not the most elegant ride, but it'll get me there (double paranthetically – not unlike my Pontiac Vibe.)) That doesn't make me a bad person, it just means my expertise lies elsewhere. I was trained as statistician and biologist, not a programmer. I can program some, just like I can change the oil on my beloved Pontiac Vibe and do other routine maintenance. The difference is that I can't pay someone to write my code for me; I can pay someone fix my car. So yeah, my code is ugly as hell. It's all zip ties, duct tape, and unnecessary if else statements that could probably vectorized if I was clever than I am. Maybe the one difference is that I'll send you my code right away if you ask.

Reply ↓

Andrew on May 8, 2020 2:19 PM at 2:19 pm said:

Dalton:

I think the counter-argument here is that, even in the absence of any government policies on social distancing, at some point people would've become so scared that they would've socially distanced themselves, canceling trips, no longer showing up to work and school, etc.

Conversely, a good argument in favor of governmentally mandated social distancing is that it coordinates expectations. I remember in early March that we had a sense that there were big things going on but we weren't sure what to do. If everyone is deciding on their own whether to go to work etc., things can be a mess. Yes, there is an argument in favor of decentralized decision making, but what do you do, for example, if schools are officially open but half the kids are too scared to show up?

Reply ↓

dhogaza on May 8, 2020 2:32 PM at 2:32 pm said:

Andrew:

"I think the counter-argument here is that, even in the absence of any government policies on social distancing, at some point people would've become so scared that they would've socially distanced themselves, canceling trips, no longer showing up to work and school, etc."

That doesn't fit well with the parallel narrative that our economic hurt has been caused by government policies, and if government hadn't done anything, economically everything would be hunky-dory. Nor by the argument (which we're seeing made here by some) that the fear is irrational and was induced by a motivated press and cadre of scientists. Implication being that without this history, there would be no fear.

I think you're being too charitable.

Reply ↓

Daniel Lakeland on May 8, 2020 2:42 PM at 2:42 pm said:

> even in the absence of any government policies on social distancing, at some point people would've become so scared that they would've socially distanced themselves, canceling trips, no longer showing up to work and school, etc.

I think it's important to point out that the projections from the model are explicitly *for the case where this isn't happening*. So let's say we project that if we continue along like worker bees that 500,000 people die.

The fact that of course we won't continue along like worker bees is not evidence that this is a wrong prediction.

Yes, we should also do some projections of "what happens if we all get scared and try to reduce our contacts and it succeeds at the level of X" which … yeah THEY DID SOME OF THIS.

https://spiral.imperial.ac.uk:8443/bitstream/10044/1/77482/14/2020-03-16-COVID19-Report-9.pdf

See figure 2 "close schools and universities" and "case isolation" and etc. The take-away from that is really "you have to do a lot in order to get significant reductions"

Reply ↓

Dalton on May 8, 2020 2:55 PM at 2:55 pm said:

I think COVID-19 does demonstrate that there is a good argument for a well-organized government response. It's the same disease in every country, yet with very different outcomes. That probably arises from both government and from culture. As you point out, the general public is going to respond in a certain way even absent governmental mandates. But different cultures would probably respond differently even given the same set of mandates. That said, there are clearly better and worse responses to this worldwide, at least when you use some combinations of metrics like per capita deaths and per capita number of new cases (and maybe decline in GDP or unemployment rate).

I hear tell of people arguing for a Swedish model for the United States (only hold the socialized medicine please). If we're all picking and choosing countries we would rather our own was like, I for one would've preferred a German model (among the lowest per capita deaths with a clear decline in the number of new cases). That said, I don't think either set of government policies would play out quite the same way in the United States. Still it'd be an interesting thought experiment to graft the government responses of a different country onto our own culture and imagine how scenario would play out.

Reply ↓

Martha (Smith) on May 8, 2020 5:07 PM at 5:07 pm said:

Good points.

Reply ↓

Twain on May 8, 2020 2:32 PM at 2:32 pm said:

Dalton & Andrew,

Something to consider: https://www.kff.org/medicaid/issue-brief/state-reporting-of-cases-and-deaths-due-to-covid-19-in-long-term-care-facilities/

LTCFs have generated ~27% of COVID-19 deaths among 23 states reporting. In many states, it is higher than ~50%.

Should we include the count of deaths from LTCFs when calculating IFR_avg (as stratified/weighted by age)?

Reply ↓

dhogaza on May 8, 2020 2:35 PM at 2:35 pm said:

If you do, please remember that the reproduction rate once a long-term facility is infected seems to be rather high than in the general population.

BTW my gf's uncle is in such a facility and covid-19 is on the move. Fortunately not in his part of it (which has been isolated, effectively, I hope).

So perhaps I'm a bit sensitive to anything that hints at a possible implication that deaths in such institutions be devalued.

Reply ↓

Dalton on May 8, 2020 2:40 PM at 2:40 pm said:

What would be the point? You'd be intentionally biasing your estimate low. Or rather by excluding long term care facilities you'd alter your scope of inference and arrive at an estimate for the age-averaged IFR that only applies "the population of X country, excluding long-term care facilities." So yeah, you'd have a somewhat lower number, but what exactly would the point?

Reply ↓

Twain on May 8, 2020 3:22 PM at 3:22 pm said:

I think excluding deaths/cases from LTCFs is necessary when determining if shutdowns are safe to end.

People at LTCFs rarely leave their facilities — their only source of mixing with the greater population is employees who work there. Therefore, residents of LTCFs are not really apart of the "general" population. Moreover, and sadly, many LTCFs have terrible conditions (under-staffing, minimal oversight, etc.) that place already vulnerable residents at higher risk for secondary-infection, delayed care, and other factors that will increase their likelihood of severe outcomes or dying — which could be skewing IFR higher than those of the same age outside of LTCFs — e.g., LTFCs is a risk factor that we should weight in addition to age when calculating IFR.

So should we consider their IFR when determining if the general population is safe to return to work and other activities?

Reply ↓

Martha (Smith) on May 8, 2020 5:24 PM at 5:24 pm said:

Twain said, "People at LTCFs rarely leave their facilities — their only source of mixing with the greater population is employees who work there. "

But the employees who work there mix with the greater population (i.e., are potential "vectors" for the spread of the disease)

Reply ↓

Daniel Lakeland on May 8, 2020 6:01 PM at 6:01 pm said:

Exactly, what "protecting LTCFs" looks like is something essentially like this:

1) Hire 3x the number of staff needed to staff the facility.

2) Group the staff into 3 groups. At all times, one group would be *housed on site* for say 3 weeks.

3) Every 3 weeks the staff would change, the incoming staff would be quarantined for the 7 days before their next shift, and then thoroughly tested prior to arrival.

4) rotate like this with staff having 3 weeks on, 5 weeks off, 1 week quarantine, test, repeat.

But you'd also have to provide *on site* housing for the active staff... and everything going in and out of the facility would need to be delivered to a quarantine storage, kept in quarantine for 5 days, and then rotated into use. This includes all medicines, all food, all equipment and supplies, etc.

This would undoubtedly raise the cost of operating a LTCF by at least a factor of 2, and they're already expensive, so it would require government intervention to make it happen.

It's not easy.

Reply ↓

Twain on May 8, 2020 7:16 PM at 7:16 pm said:

Given the enormous resources directed to-wards other NPIs and hospitals, I think it's fair to say more was possible to protect LTCFs. Mandate masks, symptom logging, checking temperature, hourly handwashing, etc., in early March to mitigate spread from worker to resident until the necessary resources became available to do more protective measures, like those you outline above.

Twain on May 8, 2020 2:44 PM at 2:44 pm said:

dhogaza,

Oh no! I'm not saying to ignore deaths in LTFCs whatsoever. All my grandparents and multiple close family friends are in LTCFs; I check on them weekly because I'm worried the next week will be when the contract the virus.

Rather, I think we should treat LTCFs as separate to the general population. And IMO, I think lumping cases/deaths in LTCFs with all others has allowed the situation to become so dire in many states.

Reply ↓

Dalton on May 8, 2020 3:19 PM at 3:19 pm said:

I still fail to see your point. What do you mean that LTCF deaths have "allowed the situation to become so dire?"

Reply ↓

Twain on May 8, 2020 3:27 PM at 3:27 pm said:

By not assuming "residence in a LTCFs" as a major risk for death along with age during the early weeks of the pandemic, many states failed to see the risk LTCFs presented and did not take early action to bolster their oversight, staffing, etc., to prevent outbreaks within them.

For example, in NJ, now seven weeks into their COVID-19 outbreak, just yesterday called the National Guard to aid LTCFs. Too little too late: COVID-19 has killed ~8% of all residents in LTCFs with more likely coming since ~40% of residents have infections (cannot find numbers on LTCFs residents hospitalized).

Reply ↓

Dalton on May 8, 2020 3:40 PM at 3:40 pm said:

Thanks Twain, I think I understand you now. I agree the situation is dire in long-term care facilities and much more should have been done to protect them in advance. Unfortunately, we haven't put a huge priority as a society in ensuring that long term care facilities are appropriately staffed, that staff are appropriately compensated and trained to ensure high quality staffing, and that facilities are well designed and well supplied. These things all cost money and the vast majority of those needing long term care are unable to afford that cost and the American electorate has tended to vote for politicians who do not put high priority of funding healthcare for those who can't afford it.

But I still don't see how that applies to estimating an age-averaged IFR. It may be the COVID-19 spreads faster in LTCF hence their representation in the death counts so far, but that doesn't mean COVID-19 is any less deadly for those in the same age demographic but whom reside at home or with family.

Reply ↓

Twain on May 8, 2020 7:12 PM at 7:12 pm said:

Dalton,

Agreed, the continued lack of support for LTCFs has been a constant frustration.

But back to the missus with IFR — using age- weighted assumes that everyone is equally susceptible. Which may be skewed toward a higher amount by the poor conditions at LTCFs that would allow the virus to build and fester.

For example, the elderly who live with family are likely to be more carefully monitored and receive care earlier, which is key to better outcomes. (I can't imagine family members disregarding a loved one…but sadly this happens at LTCFs often for the reasons you mention.) They may also be exposed to an overall higher viral load because of poor sanitation and HVAC.

Thus, it may be better to weight by "residence in LTCFs" in addition to age when calculating IFR_avg.

Reply ↓

Zhou Fang on May 8, 2020 7:32 PM at 7:32 pm said:

My understanding is that actually the Imperial College models at present, being mostly based on official death counts from Italy and China, actually implicitly ignore care home deaths because they are excluded from such counts (and are much more poorly recorded). So they potentially underestimate.

Reply ↓

Twain on May 8, 2020 7:52 PM at 7:52 pm said:

Are LTCFs much less common in China and Italy overall?

If yes, that would make LTCFs a unique risk factor for the US were they are commonplace.

Reply ↓

Twain on May 9, 2020 12:40 AM at 12:40 am said:

My prior comment was not clear. What I deserve for trying to post via my phone…

Reasons why LTCFs may make residents more _suspectible_ to poor outcomes from COVID-19:
1. Decreased probability of receiving care early because of neglect (due to understaffing, etc.).
2. Increased probability of routine exposure to higher viral load because of insufficient sanitation, HVAC, etc.
3. Increased probability of contracting a secondary-infection soon after developing symptoms for COVID-19 (LTCFs are rife with the many of the same infections that plague hospitals because they have loads of vulnerable hosts).

Based on the above factors increasing susceptibility to poor outcomes from COVID-19, I think it is necessary to treat residents of LTCFs as an independent population having its own IFR. Since ~30% of all deaths from COVID-19 in the US (23 states reporting; https://www.kff.org/medicaid/issue-brief/state-reporting-of-cases-and-deaths-due-to-covid-19-in-long-term-care-facilities/), this could skew IFR_avg for the general population higher if one does not exclude LTCF cases/deaths

from the calculation.

You could still calculate IFR_avg for use in models: have LTFC residents and everyone else be two independent groups and then calculate a weighted IFR_avg for the two groups.

I cannot check how much LTCFs are skewing IFR_avg because there aren't any US-wide (or even state-wide) age-stratified IFRs that seem reliable. If anyone has some they think are, please share! (I hesitate to use data from Italy because I'm not sure if LTCFs are similar to (or as prevalent) as those in the US.)

Reply ↓

statflash on May 9, 2020 12:46 AM at 12:46 am said:

I finally found the age and sex breakout for COVID-19 deaths (no luck with getting the same for cases) at the CDC:

https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm#AgeAndSex

I don't know how you could get infection rates by state and age group, but you could make assumptions and use the numerators from the above for calculating IFR.

statflash on May 9, 2020 12:54 AM at 12:54 am said:

Does this work for the numerators you need?

the age and sex breakout for COVID-19 deaths (no luck with getting the same for cases) at the CDC:

https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm#AgeAndSex

For the denominators (# infected) you could make assumptions and use population tables.

Twain on May 9, 2020 1:21 AM at 1:21 am said:

statflash,

Thanks for sharing! I've seen that data as well and considered using it.

I'm straying away from assuming IFR since it seems to vary in it's age-dependent skew. So if I assume wrong, then my result could be erraneous.

I did check data from NJ using their cases and deaths (same age-grouping). If you exclude deaths from LTFCs, it causes a ~3x decrease in CFR_avg (~6% to 2%). So including LTCFs as a clear effect. But without knowing true number of infections, I can't determine if this effect is big, small, or negligible.

Carlos Ungil on May 9, 2020 4:35 AM at 4:35 am said:

> this could skew IFR_avg for the general population higher if one does not exclude LTCF cases/deaths from the calculation.

But they are part of the general population!

It's true that even at the population level the IFR is not a well defined because the population may be "all the people", "people who has been infected so far", or "people infected in step 25 of run #42 of a badly-written simulation" and in those actual or hypothetical cases the number of fatalities depends on many factors. If we are interested in the first version (because all are going to be infected) looking at the second one (deaths in those infected so far) may result in an underestimate or an overestimate, depending on the relative prevalence of the infection and underreporting of infections and deaths in different subgroups.

Twain on May 9, 2020 9:35 AM at 9:35 am said:

Carlos,

> But they are part of the general population!

Based on Daniel's above explanation and my further reading, it seems the major assumptions when calculating population-weight IFR_avg are (1) all ages are equally susceptible and (2) all ages have no outside factors that could skew the probability of severe outcomes.

As I outline above: (1) does not hold for LTCFs because the likelihood of higher viral exposure increases; (2) does not hold because likelihood of neglect, deprivation, secondary-infection are much higher in LTCFs than elsewhere, which lead to late/delayed treatment. Therefore, residents of LTCFs have specific factors making them more susceptible to infection AND severe outcomes.

Therefore, I think it is best to treat those in LTCFs and the rest of the population separately when calculating IFR_avg and subsequently applying it in models. If this does not make sense, please explain why — I've definitely be mistaken before!

Martha (Smith) on May 9, 2020 12:36 PM at 12:36 pm said:

Twain said,
"Therefore, I think it is best to treat those in LTCFs and the rest of the population separately when calculating IFR_avg and subsequently applying it in models. If this does not make sense, please explain why — I've definitely be mistaken before!"

I'm not convinced that "it is best to treat those in LTCF's and the rest of the population separately when calculating IFR_avg". I can't help but wonder if perhaps people are stuck on a certain kind of model and not considering that they need to "think outside of that box". Here's what I see as problematical:

First, I think it would be better to consider the entire category of "congregate living" as the appropriate grouping, rather than just focusing on LTCF's. In other word, instead of a model that considers the two groups "people who live in LTCF's" and "people who live in family or similar small units", I think that a good model needs to consider the two groups "people in congregate living" and "people who live in family or similar small units" By "congregate living" I mean "people who live in large group residences such as LTCF's, orphanages, residential schools, convents, monasteries, prisons, …".

My reasoning is, first, that separating out just LTCF's produces a confounding between age and type of living unit .

Second, the virus does not confine itself to spreading between people living in the same living unit (or type of unit). As I have mentioned in a previous post, there is interaction (and likely spread) of the virus between people in LTCF's and people living in small units.

Someone (I forget who) responded to that previous post by suggesting it made things too complex. Yes, it does make for a more complex model. I agree that "all models are wrong, some are useful", but I think that the current models are too wrong to be useful.

Martha (Smith) on May 9, 2020 12:41 PM at 12:41 pm said:

I just wrote, "Someone (I forget who) responded to that previous post by suggesting it made things too complex."

I now see that it was a comment by Zhou Fang yesterday that appears further down in the thread.

Joshua on May 9, 2020 12:45 PM at 12:45 pm said:

Martha –

> First, I think it would be better to consider the entire category of "congregate living" as the appropriate grouping, rather than just focusing on LTCF's.

Just a thought – that might be more relevant to infection rate as opposed to fatality rate?

Relatedly, I just posted this: https://statmodeling.stat.columbia.edu/2020/05/08/so-the-real-scandal-is-why-did-anyone-ever-listen-to-this-guy/#comment-1333409

My guess is that meat packers would be far at one end of the spectrum in terms of "congregate living" but with practically no people over 60. Could be an interesting sample for comparison.

Martha (Smith) on May 9, 2020 3:49 PM at 3:49 pm said:

Joshua said,

"Martha –

> First, I think it would be better to consider the entire category of "congregate living" as the appropriate grouping, rather than just focusing on LTCF's.

Just a thought – that might be more relevant to infection rate as opposed to fatality rate?"

It's hard to say — partly since higher infection rate might lead to higher fatality rate (e.g., by increasing viral load, which is often mentioned as increasing fatality rate.)

Joshua also said,
"My guess is that meat packers would be far at one end of the spectrum in terms of "congregate living" but with practically no people over 60. Could be an interesting sample for comparison."

Good point. So I think what we really need to think about is "high chance of exposure" and/or "high risk of repeated exposure (i.e., of high viral load)". This would include congregate living, but also what (for lack of a better term) I will call "high exposure working" — which includes meat packers, and also people working in nursing homes, and probably other occupations as well. The crux of the matter really is to consider "age" and "high exposure environments" as separate variables, whereas just including or excluding LTCF's confounds these two variables, and doesn't include all "high exposure" possibilities.

Joshua on May 9, 2020 6:47 PM at 6:47 pm said:

Martha –

> The crux of the matter really is to consider "age" and "high exposure environments" as separate variables,

Yes, I was just thinking about that this morning. Especially since there is now so much focus in the media on number of deaths in LTCFs.

Twain on May 9, 2020 10:39 PM at 10:39 pm said:

Martha,

> I'm not convinced that "it is best to treat those in LTCF's and the rest of the population separately when calculating IFR_avg".

Yikes! I was not implying that separating LTCF residents from everyone was *only* way stratify when calculating IFR_avg. Other factors are possible — as you mention!

> I can't help but wonder if perhaps people are stuck on a certain kind of model and not considering that they need to "think outside of that box"

Agreed. Something that surprises me is how homogenous these models treat the population with regards to infection, susceptibility, outcomes, mortality, etc. People are VERY heterogeneous in how the respond to disease, even within specific ranges of age. We see this for many diseases, like cancer, influenza, streptococcus pneumoniae, etc.

We can account for some of this immunity by weighting by age, for example; but we first have to determine/recognize these factors and how they stratify across a population.

For example, how does innate immunity affect the susceptibility of of those age 0-30 to develop an infection capable of transmitting to others? If innate immunity allows them to become infected but not infectious (for most cases), that has huge bearing on how we model and combat the virus.

> First, I think it would be better to consider the entire category of “congregate living” as the appropriate grouping, rather than just focusing on LTCF’s. […] Second, the virus does not confine itself to spreading between people living in the same living unit (or type of unit). As I have mentioned in a previous post, there is interaction (and likely spread) of the virus between people in LTCF’s and people living in small units.

An excellent point! And I agree. It is clear congregate living confers enough differences in susceptibility (to infection and poor outcomes) that one cannot consider them similar to those living the more traditional “familial style” living.

> I agree that “all models are wrong, some are useful”, but I think that the current models are too wrong to be useful.

Agreed. Maybe the current simplifications will work. But maybe they won’t. We have to check given the clear skewing of hospitalizations, deaths, etc., for certain deprived, at-risk, and elderly populations compared to everyone else.

> The crux of the matter really is to consider “age” and “high exposure environments” as separate variables,

Excellent point (!) and exactly what I concluded after reading your posts. Well said overall.

—

Joshua,

Thank you for sharing! I’ll check the data and see what I find.

As we’ve discussed, a high-quality statewide or nation-wide study of seroprevalence cannot come soon enough!

Martha (Smith) on May 8, 2020 5:21 PM at 5:21 pm said:

Twain asks, “Should we include the count of deaths from LTCFs when calculating IFR_avg (as stratified/weighted by age)?”

I’m not clear just what you’re asking. Are they not currently included? Or are you asking if maybe they should be excluded?

But maybe instead, the method of calculating IFR_avg meeds to be stratified/weighted according to two variables: age, and residing in a LTCF? This makes sense to me, since (as you point out) the proportion of deaths of LTCF residents from COVID-19 can run very high. In other words, it sounds like “lives in an LTCF” is an important variable to take into account in studying the epidemic.

Reply ↓

Martha (Smith) on May 8, 2020 5:27 PM at 5:27 pm said:

My apologies for not having read earlier comments (some of which have addressed my comments) before I posted this one.

Reply ↓

Zhou Fang on May 8, 2020 5:27 PM at 5:27 pm said:

Well, you can always make a model arbitrarily more complicated to capture reality. The trade off is between the improvement you get in terms of reducing approximations, and the error you open yourself up to in terms of potentially mis-estimating parameters, getting erroneous data, having to contend with incomplete data, making incorrect assumptions, coding errors…

And also all that effort……. :P

Reply ↓

yyw on May 8, 2020 2:59 PM at 2:59 pm said:

Ferguson model predicted 6.4k-20k UK deaths over 2 years assuming school closure, social distancing, and case isolation for 5 months. UK deaths stand at 31k not even 2 months into lockdown. It's clear that the model prediction was too low. The report did not provide a projected trajectory of fatalities, otherwise it can be compared with actual trajectory to see how far off it has been.

It could be because actual social distancing is lower than he assumed, but his assumptions of 25% universities remaining open, 50% less work contact, 75% less other social contact, and 75% compliance did not seem too stringent and may actually be less stringent than reality.

Reply ↓

Dalton on May 8, 2020 3:07 PM at 3:07 pm said:

yyw,

Are we talking about the same report? I was referencing this one from March 16th : https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf

The projected trajectory of fatalities is in Figure 1.

Reply ↓

yyw on May 8, 2020 3:35 PM at 3:35 pm said:

That's for the unmitigated scenario only. I am more interested in trajectories under different policy assumptions.

Reply ↓

Zhou Fang on May 8, 2020 3:43 PM at 3:43 pm said:

If I recall correctly, his team announced that the initial model used an estimate of R0 that was subsequently thought to be too small (due to the role of asymptomatics not being well understood at that time.)

Reply ↓

David Young on May 10, 2020 12:00 AM at 12:00 am said:

Dalton, Ferguson's IFR was not unreasonable and agreed with Fauci's guess. However, it was nothing more than a guess.,

1. In all past epidemics IFR's have always declined over time, sometimes by orders of magnitude. Even a smart layman would take that into account.
2. There are by now at least 10 serologic data sets around the world. They pretty much uniformly show an IFR less than 0.5% with the best ones showing perhaps 0.12% to 0.31%. These studies show that the virus has infected vastly more people that the worthless "case" statistics trumpeted by irresponsible journalists.
3. Total fatality numbers are misleading and would not be used by competent scientists. The number that is useful is the total excess mortality number. By Ferguson's own estimation, 2/3 of those who die with covid19 would have died within the year because they were already seriously ill. That's borne out by the high percentage of fatalities among nursing home residents. Thus for policy making, 1 million becomes 333,000, about 13% excess mortality.

Bottom line Ferguson missed by quite a bit and on the alarmist side, which seems to be his track record.

Reply ↓

dhogaza on May 10, 2020 12:43 AM at 12:43 am said:

"Bottom line Ferguson missed by quite a bit and on the alarmist side"

Climate science denialist trots out the "alarmist" word.

Perhaps he's unaware of the criticisms of the "the best [study] showing perhaps 0.12% to 0.31%" made by our host, a noted statistician?

But given data that shows an R0 value of about 3, even with an IFR as low as 0.31% you're talking 667K dead in the US without mitigating factors. More than 10 times the number that died in the worst pandemic flu season in the US this past decade.

Why aren't you over at James Annan criticizing his model? Surely an expert like you can rip it to shreds.

Reply ↓

David Young on May 10, 2020 10:22 AM at 10:22 am said:

I'm not denying anything. Name calling is for children.

I was actually referring to the Miami Dade data.

It is misleading to give total fatalities. The standard to use is excess mortality. If 2/3 of those who die with Covid19 would have died within a year, the excess mortality would be 222K or about 8% on an annual basis.

The problem with modeling is that's its an ill-posed problem because of the exponential growth. Results are sensitive to parameters. The details of transmission and individual behavior have a large effect. I doubt all models in this setting. The best we can do is random testing to get a handle on some of the details.

Reply ↓

Martha (Smith) on May 10, 2020 4:23 PM at 4:23 pm said:

David said,
"The problem with modeling is that's its an ill-posed problem because of the exponential growth. Results are sensitive to parameters. The details of transmission and individual behavior have a large effect. I doubt all models in this setting. The best we can do is random testing to get a handle on some of the details."

Definitely points worth taking into account when using models.

Reply ↓

Twain on May 10, 2020 5:49 PM at 5:49 pm said:

> The problem with modeling is that's its an ill-posed problem because of the exponential growth. Results are sensitive to parameters. The details of transmission and individual behavior have a large effect.

Agreed. This is *especially* a problem if factors independent of age — socioeconomic status, occupation, innate immunity, mobility, location — affect susceptibility by increasing routine exposure to virus, increasing probability of high initial exposure to virus, or increasing severe outcomes to virus. As time progresses, this seems to be the case, as very high IFR among LTCFs and congregate housing seems to indicate.

Reply ↓

Stephen M St Onge on May 12, 2020 1:05 AM at 1:05 am said:

David Young:

"If 2/3 of those who die with Covid19 would have died within a year," and the key word is "if."

What reason is there to believe any of this?

Reply ↓

dhogaza on May 10, 2020 4:43 PM at 4:43 pm said:

"Name calling is for children."

After trotting out the predictable "alarmist" word.

"If 2/3 of those who die with Covid19 would have died within a year"

Which, of course, despite the fact that most who die are older, diabetic, etc is not close to true.

Reply ↓

Twain on May 10, 2020 6:28 PM at 6:28 pm said:

> Which, of course, despite the fact that most who die are older, diabetic, etc is not close to true.

Can you provide source(s) for this?

– Data from NYC shows 0.6% of deaths had no underlying conditions and 53% of deaths were age 65+ (https://www1.nyc.gov/site/doh/covid/covid-19-data.page; note NYC uses age 45-64 grouping; how many of this group at age 60+?).
– Data from NJ shows ~79% of deaths were age 65+ (https://www.nj.gov/health/cd/topics/covid2019_dashboard.shtml; no data on underlying conditions).
– Data from Italy shows 4% of deaths had no underlying conditions and 94% of deaths were age 60+ (https://www.epicentro.iss.it/).
– Data from Spain indicates that 95% of deaths were age 60+ and 95% of deaths had underlying conditions (https://www.mscbs.gob.es/profesionales/saludPublica/ccayes/alertasActual/nCov-China/situacionActual.htm; see pg 11, table 7 of "Epidemiological analysis COVID-19").

The data seem to make clear that most who date are older and/or having underlying conditions compose most deaths from COVID-19.

Reply ↓

Carlos Ungil on May 10, 2020 6:59 PM at 6:59 pm said:

You may have misread his comment. He explicitly acknowledges that most who die are older, diabetic, etc. and your data doesn't show that most of them would have died within a year.

Reply ↓

dhogaza on May 10, 2020 7:44 PM at 7:44 pm said:

Carlos Ungil

Bingo!

Twain on May 10, 2020 7:55 PM at 7:55 pm said:

Looks like I did; sorry about that!

Stephen M St Onge on May 12, 2020 1:15 AM at 1:15 am said:

Twain:
"Data from NYC shows 0.6% of deaths had no underlying conditions and 53% of deaths were age 65+."

Last time I looked, it showed no such thing. It showed 0.6% _definitely_ had no known comorbidities, and 23% or so there was no information. So the IFR based on that varies by a potential factor of forty for no known comorbidities. "Absence of evidence is not evidence of absence."

And when the comorbidities include obesity, asthma, and diabetes, the proportion of the population that have at least one may be over half the total population of the U.S.

Reply ↓

Twain on May 12, 2020 8:44 AM at 8:44 am said:

Stephen,

My prior is this: I have been following the data from NYC daily since its release. During that time, I have seen almost all of the "unknown" category eventually distributed to the "underlying conditions" category — which has held the "no underlying conditions" category steady at ~2-0.5% for over 3 weeks. And if we are really going to nit-pick — during that time I have seen the count of "no underlying conditions" drop from 152 on April 16 to 42 on April 17; so if anything, NYC may have been over-counting those in the category.

Of course, I don't know for sure if it is really 0.6% having no underlying conditions, but I think the above trends make that more likely than otherwise.

David Young on May 11, 2020 11:58 AM at 11:58 am said:

Dhog, I would suggest you actually look at some science on this matter of who is dying from covid19. Even Ferguson's report shows the exceptionally strong age dependence. Everyone agrees that this disease is perhaps 1000 times more deadly among those over 80 than in those under 20. There is a uniform increase with age. There is a gradual increase with age of conditions such as diabetes and the metabolic syndrome. Most people will eventually become diabetic by the standard criteria. They have recently increased the A1C threshold for people over 80 to 7.1%. Ioaniddis and some colleagues have a paper analyzing the Italian data that is quite good. So, the 2/3 number is Ferguson's estimate and I think is perhaps a lower bound.

Reply ↓

Zhou Fang on May 11, 2020 12:07 PM at 12:07 pm said:

https://wellcomeopenresearch.org/articles/5-75

" Results: Using the standard WHO life tables, YLL per COVID-19 death was 14 for men and 12 for women. After adjustment for number and type of LTCs, the mean YLL was slightly lower, but remained high (13 and 11 years for men and women, respectively). The number and type of LTCs led to wide variability in the estimated YLL at a given age (e.g. at ≥80 years, YLL was >10 years for people with 0 LTCs, and <3 years for people with ≥6). "

So, no.

Reply ↓

dhogaza on May 11, 2020 1:37 PM at 1:37 pm said:

Zhou Fang:

The last thing to expect from DK Young is honesty.

I also saw a preprint that suggested that the YLL value for seasonal flu is actually worse than for covid-19. I didn't think to save a link to it …

Oh, and his strange contraction of "dho gaza" to "Dhog" was begun by an ex-climate science denier years ago. David's been around for years playing the climate science denialist game.

confused on May 11, 2020 2:53 PM at 2:53 pm said:

Hmmm. I have to say it seems really odd that adjusting for conditions would only change the result by one year – less than 10%!

Perhaps this is because nearly all elderly people have some pre-existing conditions? But what about differences in severity? The paper does say they didn't have data on severity of underlying conditions. This could make a big difference, as some of the conditions mentioned (especially hypertension and diabetes) are very common and people live decades with them if well-controlled.

I never believed the idea that the years-lost would be *really* small (ie usually <1 year) as some have suggested, but this does seem like a surprisingly high value given how large a proportion of the deaths are in nursing homes/LTCFs (who are presumably less healthy than average even among their age group).

David Young on May 11, 2020 4:10 PM at 4:10 pm said:

I glanced at your reference Zhou it appears to be based on a model.

Data in the US shows that perhaps 40%-60% of fatalities are amoung nursing home residents. Those in nursing homes generally have a low life expectancy and are often seriously ill.

The 2/3 number is Ferguson's number. You can also look at the Ioannidis paper on analyzing the Italian data. That's more believable to me because its data based.

dhogaza on May 11, 2020 4:27 PM at 4:27 pm said:

In NY it's about 20% dead in long term care facilities. US government says that the average stay in a long term care facility is > 3 years for women, 2.2 years for men. YLL > 1 when they die from covid-19.

Note: not all patients long term care facilities are in nursing homes.

Zhou Fang on May 11, 2020 5:01 PM at 5:01 pm said:

Where exactly did Ferguson say that 2/3rds of those who died would have died within the year? If he did make such a ridiculous claim that would cost him a fair bit of credibility with me.

David Young on May 11, 2020 5:27 PM at 5:27 pm said:

I didn't save the link. It was in a report where Ferguson was calculating excess morality.

In King County Washington, over 50% of fatalities have been in residents of nursing homes. Of course, the majority of deaths occurred in a hospital. Some of these were infected in hospital and were already seriously ill with something else.

Zhou Fang on May 11, 2020 5:54 PM at 5:54 pm said:

Oh I found it:

> "We don't know what the level of excess deaths will be in this epidemic," Ferguson said. In other words, we don't know the extent to which COVID-19 will increase annual deaths above the level that otherwise would have been expected. "By the end of the year, what proportion of those people who've died from COVID-19 would have died anyhow?" Ferguson asked. "It might be as much as half to two-thirds of the deaths we're seeing from COVID-19, because it's affecting people who are either at the end of their lives or in poor health conditions. So I think these considerations are very valid."

So "We don't know [...] it might be as much as" gets turned into "By Ferguson's own estimation, 2/3 of those who die with covid19 would have died within the year because they were already seriously ill."

I'm not going to pretend to be surprised here, this is a pretty repeated pattern of deliberate deception coming from you.

Joshua on May 11, 2020 6:01 PM at 6:01 pm said:

Zhou –

Don't assume malignant intent when sheer incompetence will suffice as an explanation.

Zhou Fang on May 11, 2020 6:06 PM at 6:06 pm said:

Joshua:

I think if the error distribution is sufficiently asymmetric, then I don't think simply high levels of variability is good enough as an explanation.

David Young on May 11, 2020 6:16 PM at 6:16

pm said:

well yes its an estimate. If you are really interested in this, you might try combing through the Ioannidis et al paper analyzing the Italian data. If you only want to find some small thing I said that can be questioned, I give you the bronze medal for nit picking. My memory is vague but I think I heard this 2/3 number somewhere else perhaps in Ioannidis' video.

Joshua on May 11, 2020 6:17 PM at 6:17 pm said:

Zhou –

Add heavily motivated reasoning to incompetence. It's a powerful coupling.

Some people have an established history of saying things that are just obviously and embarrasingly wrong, such that no one serous could consider them to be remotely plausible. In such a situation malignant intent becomes less likely.

David Young on May 11, 2020 6:23 PM at 6:23 pm said:

Do either Zhou or Josh have any substantive science on the subject of excess mortality? No. It's pure mind reading and childish speculation.

Zhou Fang on May 11, 2020 6:43 PM at 6:43 pm said:

Joshua:

I think that sort of motivated unapologetic incompetence becomes indistinguishable from malice.

David:

Do *you* have any substantive science? Because right now the entire basis of your argument is shown to be based on what might charitably be considered a case of misremembering. If you want to make this case why should other people dig through these papers for you?

Also yes if Ioannidis is claiming this that discredits him also.

David Young on May 11, 2020 7:23 PM at 7:23 pm said:

Zhou, You are missing the main point. Excess mortality is a very minor part of my point.
The main point is that serological studies are now pretty common and at least 4 of them show IFR's between 0.12% and 0.41%. If you have any information to add, I'd entertain it, but I'm not going to engage your essentially ad hominem attacks.

Zhou Fang on May 12, 2020 7:30 AM at 7:30 am said:

And when we dig in and find that statement is also a misrepresentation, then this too will become a "very minor part of my point".

Andrew on **May 12, 2020 8:11 AM at 8:11 am** said:

David:

When people point out false or misleading things you've said, it's inappropriate to give them "the bronze medal for nit picking." This is a statistics blog. Details matter, and if you don't care about details, you're kinda wasting our time here.

If people go to the trouble of carefully reading what you've written and they find untrue or misleading statements, you should thank them for their effort, not criticize them for "nit picking." If you can't appreciate when people correct you, you should avoid open discussion forums like this blog. It would make more sense for you to be giving Ted talks or writing articles for PNAS or Perspectives on Psychological Science.

David Young on **May 12, 2020 9:39 AM at 9:39 am** said:

Andrew, I replied but it ended up below in the wrong place below.

Willard on **May 12, 2020 10:15 AM at 10:15 am** said:

Andrew,

To "you're nitpicking" we could add:

> *Well, You are quoting Gelman out of context.*

https://andthentheresphysics.wordpress.com/2020/05/09/attacking-scientists-who-tell-the-truth/#comment-175738

That was in response to me quoting your own summary.

There is also:

> *It's odd that people here are focusing on criticisms of inadequate data and method info in the paper while ignoring the fact that the criticisms are from a person who is not an expert and admits that the authors are the experts.*

https://andthentheresphysics.wordpress.com/2020/05/09/attacking-scientists-who-tell-the-truth/#comment-175821

Willard on **May 12, 2020 10:35 AM at 10:35 am** said:

Andrew,

Seems that pressing enter sends the post. I meant to add this other gem from James about your favorite aeronautic engineer:

> *I expect David will appear shortly and acknowledge that he was wrong and that Nic massively underestimated the death rate and that in fact mortality across much of Europe is at an extremely high level despite strong attempts to suppress the epidemic. I can see his flying pig coming in to land right now.*

http://julesandjames.blogspot.com/2020/04/euromomo_10.html

You should check James' model.

David Young on **May 12, 2020 1:04 PM at 1:04 pm** said:

Some facts about nursing homes in the US from city journal. 1.3 million residents. 39% are 85 years old or older. 19% of total US deaths. That's about 451K deaths per year or about 35% of all residents. Nursing-home residents are extremely vulnerable to the coronavirus due to multiple comorbidities: 72 percent have hypertension, 38 percent heart disease, and 32 percent diabetes. In 2016, 45 percent of nursing homes had infection-control deficiencies; as many as 380,000 people die every year of infections in long-term care facilities.

In many juristictions more than 50% of covid19 deaths have been among nursing home residents. I think its safe to say that most of those were among those with the most serious comorbidities. That suggests that at least in this settings, many of those fatalities would have occured with a year or two regardless and that covid19 is just one of many infections that hasten those deaths. So yes, there are a lot of people who are on death's door. They are in settings where they are very vulnerable and deserve much more attention than they are getting.

dhogaza on **May 12, 2020 1:17 PM at 1:17 pm** said:

DK Young

"39% are 85 years old or older. 19% of total US deaths. That's about 451K deaths per year or about 35% of all residents"

Is this meant to support your claim that the average YLL is less than one?

dhogaza on **May 12, 2020 1:24 PM at 1:24 pm** said:

DK Young

"In many juristictions more than 50% of covid19 deaths have been among nursing home residents."

Well, sure, and I've spent a lot of time in a rural county in SE Oregon where cows outnumber people by 10:1. I'm not going to extrapolate that ratio to the country at large, though.

Joshua on **May 12, 2020 1:30 PM at 1:30 pm** said:

dhogaza –

> I'm not going to extrapolate that ratio to the country at large, though.

Over multiple online forums, over a period of weeks, our friend has demonstrated some significant difficulty with applying the concept of representative sampling.

Atomsk's Sanakan on **August 24, 2020 6:12 AM at 6:12 am** said:

@Andrew

Re: "When people point out false or misleading things you've said, it's inappropriate to give them "the bronze medal for nit picking." This is a statistics blog. Details matter, and if you don't care about details, you're kinda wasting our time here.
If people go to the trouble of carefully reading what you've written and they find untrue or misleading statements, you should thank them for their effort, not criticize them for "nit picking.""

Well-said. So now that a few months have passed, we have more evidence that addresses untrue and/or misleading statements David Young made. Though given his track record, I predict he won't admit he was wrong, as he's avoiding doing when downplaying the risk of human-induced climate change just as he illegitimately downplayed the risk of SARS-CoV-2-induced COVID-19:
http://www.realclimate.org/index.php/archives/2019/12/how-good-have-climate-models-been-at-truly-predicting-the-future/comment-page-3/#comment-752395

@David Young

Re: "Even using Ferguson's IFR's which are probably substantially too high."
https://statmodeling.stat.columbia.edu/2020/05/08/so-the-real-scandal-is-why-did-anyone-ever-listen-to-this-guy/#comment-1336470

Re: "Using Ferguson's IFR numbers by age cohort (which are probably at least a factor of 2 too high) [...]."
https://judithcurry.com/2020/05/06/covid-discussion-thread-vi/#comment-916993

Ferguson et al.'s IFR was ~0.9% for Great Britain:
https://spiral.imperial.ac.uk:8443/bitstream/10044/1/77482/14/2020-03-16-COVID19-Report-9.pdf

That did not over-estimate observed seroprevalence-based IFR, including if one takes age-stratification into account:
https://www.medrxiv.org/content/10.1101/2020.07.23.20160895v3
https://www.medrxiv.org/content/10.1101/2020.08.12.20173690v2 [see supplementary table S2(a)]
https://doi.org/10.1016/S2468-2667(20)30135-3
https://www.medrxiv.org/content/10.1101/2020.05.03.20089854v4

Re: "There are at least now 10 more meaningful serological studies from around the world.
1. There is a Danish one of blood donors Joshua pointed out. IFR is 0.08.
2. There is the Santa Clara study which was strengthened by a revision.
3. There is a Los Angeles County study which shows a low IFR too.

4. Miami Dade county which shows an IFR of 0.17-0.31% even when I took fatalities from 21 days after the mean testing date.
5. State of Arizona comes in around 0.28%.
[...]
I personally think the US as a whole is more similar to the 4 US datasets I mentioned above."
https://judithcurry.com/2020/05/06/covid-discussion-thread-vi/#comment-916993

Re: "There are by now at least 10 serologic data sets around the world. They pretty much uniformly show an IFR less than 0.5% with the best ones showing perhaps 0.12% to 0.31%."
https://statmodeling.stat.columbia.edu/2020/05/08/so-the-real-scandal-is-why-did-anyone-ever-listen-to-this-guy/#comment-1333886

Blood donor studies under-estimate IFR since, for example, they exclude very young people and very old people, instead focusing on an age-range in which people are more likely to go out and interact with people, thereby being more likely to get infected:
https://www.medrxiv.org/content/10.1101/2020.07.23.20160895v3

Also, if you apply blood donor studies to matching age-groups in the general population in order to calculate IFR, you'll leave out older people in the general population. That not only over-estimates infection rate but also under-estimates deaths, since older people are more likely to die of COVID-19. So it's on par with saying 'breast cancer is not that deadly', by performing a study that disproportionately excludes older women. Example of this under-estimation:
https://academic.oup.com/cid/advance-article/doi/10.1093/cid/ciaa849/5862661

The Santa Clara suffered from a number of crucial flaws, including a recruitment design that would over-estimate seroprevalence, reportedly misleading statements used to recruit subjects, inadequate corrections for sensitivity/specificity, funding issues, etc.
https://academic.oup.com/jid/article/doi/10.1093/infdis/jiaa429/5872489
https://www.medrxiv.org/content/10.1101/2020.04.24.20078824v1
https://www.medrxiv.org/content/10.1101/2020.05.03.20089201v1
https://tandfonline.com/doi/full/10.1080/13669877.2020.1778771
https://www.nature.com/articles/d41586-020-01095-0
https://science.sciencemag.org/content/368/6489/350.full
https://buzzfeednews.com/article/stephaniemlee/stanford-coronavirus-neeleman-ioannidis-whistleblower
https://buzzfeednews.com/article/stephaniemlee/stanford-coronavirus-study-bhattacharya-email

It's very unlikely to be representative of the general population in Santa Clara. Fortunately, there's a better study underway in Santa Clara: http://med.stanford.edu/epidemiology-dept/research/covid-research-collaborative/CA-Facts.html

The Los Angeles County study was followed by a large increase in deaths, making deaths hard to pair with corresponding infections and thus reducing the reliability of IFR inferred from it: https://www.medrxiv.org/content/10.1101/2020.07.23.20160895v3

Miami-Dade later decreased their seroprevalence estimate from 6% (thereby increasing any IFR estimated from it), possibly due to issues with their original testing protocol:
old: https://miamidade.gov/releases/2020-04-24-sample-testing-results.asp
new: https://web.archive.org/web/20200727041934/http://www.sparkc.info/

There's no evidence I know of that the Arizona results were from a randomized sample representative of the general population in Arizona: https://twitter.com/GidMK/status/1288963702905958401

We now also have other studies from different parts of the US showing an IFR of 0.5% or more. And if we look more globally at seroprevalence-based IFR estimates, the studies with the least risk of bias show an IFR of ~1%:
https://twitter.com/GidMK/status/1285020775892709377 [https://www.medrxiv.org/content/10.1101/2020.05.03.20089854v4] https://www.medrxiv.org/content/10.1101/2020.07.23.20160895v3

Some other examples of seroprevalence-based IFR estimates from the USA:
1.6% : https://wwwnc.cdc.gov/eid/article/26/11/20-3029_article
1.3% : https://louisville.edu/medicine/news/phase-ii-results-of-co-immunity-project-show-higher-than-expected-rates-of-exposure-to-novel-coronavirus-in-jefferson-county
0.8% (higher if one accounts for right-censoring) : https://washoecounty.us/outreach/2020/07/2020-07-08-jic-update-0708.php
0.6% : https://sciencedirect.com/science/article/pii/S1047279720302015
0.6% : https://cdc.gov/mmwr/volumes/69/wr/mm6929e1.htm
consistent with 0.5% – 1.0% : http://archive.is/JXtUt#selection-1159.0-1163.200 [https://rivcoph.org/Portals/0/Documents/CoronaVirus/July/News/7.27.20%20antibody%20testing%20results.pdf?ver=2020-07-27-144931-703&timestamp=1595886602504]

Hence why the CDC upgraded their best estimate of IFR to 0.65%, even though that's also likely an under-estimate, as acknowledged by a soon-to-be-updated draft of the study they cite for that estimate: table 1 of: https://web.archive.org/web/20200823001414/https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html

"Restricting the analysis to only those studies at a low risk of bias resulted in modestly reduced heterogeneity and an increased IFR of 0.76% (0.37-1.15%). [...] It is not unlikely that, after correcting for excess mortality not captured in official death reporting systems, the IFR of COVID-19 in most populations would be higher than 1%."
https://www.medrxiv.org/content/10.1101/2020.05.03.20089854v4.full.pdf

Update:
https://twitter.com/GidMK/status/1285020775892709377

David Young on May 11, 2020 2:51 PM at 2:51 pm said:

Yes, dhogGaza name calling is for children and you.

I have been right about climate models from day one as Palmer and Stevens recent paper confirms. Your contributions have always been pretty much limited to name calling and casting doubt on the settled science of turbulent flow simulations. Mind reading people's positions was also a strength of yours and nothing has changed. So far you have offered nothing meaningful here either.

Ioannidis and his scores of collaborators are honest scientists trying to find the truth. Do you agree?

Zhou, what is the yll?

Reply ↓

dhogaza on May 11, 2020 3:34 PM at 3:34 pm said:

DK Young

"I have been right about climate models from day one as Palmer and Stevens recent paper confirms.

You mean the paper that describes your behavior perfectly?

"This leaves the scientific consensus on climate change vulnerable to specious arguments that prey on obvious model deficiencies"

dhogaza on May 11, 2020 3:37 PM at 3:37 pm said:

Andrew:

I apologize for drifting far afield here.

David:

yll is years of life lost. Zhou's making it clear that it is not < 1.

David Young on May 11, 2020 6:28 PM at 6:28 pm said:

DhoGaza, What a silly meaningless attack. On the scientific substance, I am and was right. And every competent CFD expert knows this too. How about "circling the wagons" as a correct description of the communities response (and your response) to valid criticism. That is not a specious argument.

dhogaza on May 11, 2020 8:25 PM at 8:25 pm said:

DK Young

This is mostly for those who've not run into you in the past ...

Palmer and Stevens, after pointing out that basic thermodynamics etc are enough to demonstrate that we're facing real problems, say that ...

"Comprehensive climate models have been effective and essential to address the concern that such a basic understanding could be overly simplistic (i.e., missing something important, such as the existence of a mode of internal variability, which could, if it were to exist, explain trends in global mean temperature). "

In other words, existing climate models are sufficient to shut down one common denialist argument.

"When it comes to global climate change, it is what the present generation of comprehensive climate models do not show—namely, a sensitivity of global changes to either the vagaries of unpredictable regional or global circulations or effects of processes neglected in simpler models—which makes them such a powerful confirmation of inferences from basic physics and simple models."

And their weakness regarding projecting down at smaller scale doesn't change the Big Picture.

The paper is full of criticisms of the current generation of GCMs but let us not forget that they basically endorse them as being sufficient to tell us that climate change due increased CO2 is real, as are the problems for humanity that are implied by this fact.

Their concerns focus on the relatively poor performance of GCMs at the regional scale to provide the kind of projections needed for planning at that scale.

"And every competent CFD expert knows this too."

As it turns out Palmer and Stevens have a solution for the current state of affairs! It involves much greater resolution using far greater computer power, made available by paying for far greater resources than being used today, so that more finer-scale processes can be modeled by computational fluid dynamics, and fewer finer-scale processes be parameterized.

I'm not sure they're on your side, DK Young.

"What is needed is the urgency of the space race aimed, not at the Moon or Mars, but rather toward harnessing the promise of exascale supercomputing to reliably simulate Earth's regional climate (and associated extremes) globally."

"We are suggesting a new approach to climate model development (23). This approach should aim to reduce climate models' dependence on subgrid parameterizations where possible and account for their uncertainty where not. To be successful, this approach must master and motivate technological innovations, particularly in computing, and be given a sense of purpose commensurate to the task at hand."

"Global storm and ocean-eddy resolving [O(1 km)] models make it possible to directly simulate deep convection, ocean mesoscale eddies, and important land–atmosphere interactions. Prototypes of such models are already being developed (21), 3 examples of which are compared with a satellite image. By avoiding the need to represent essential processes by semiempirical parameterizations, the simulated climate of such a model is more constrained by the laws of physics."

Directly simulating deep convection, ocean mesoscale eddies, and important land-atmosphere interactions sounds like problems in computational fluid dynamics to me … doesn't sound to me that they buy into your "CFD is the achilles heel of climate modeling" argument to me.

Or your claims that climate modelers are incompetent, are stuck in the 1950s, blah blah blah.

Quick, run over there and tell them they're wrong.

dhogaza on May 11, 2020 8:38 PM at 8:38 pm said:

The last, I promise …

DK Young:

"How about "circling the wagons" as a correct description of the communities response (and your response) to valid criticism."

What they actually say is …

"In our view, the political situation, whereby some influential people and institutions misrepresent doubt about anything to insinuate doubt about everything, certainly contributes to a reluctance to be too openly critical of our models. Unfortunately, circling the wagons…"

A bit more nuanced …

Of course, the professionals you've crossed swords don't accept most of your criticisms as valid so in your case the point is moot …

David Young on May 11, 2020 9:59 PM at 9:59 pm said:

DhogGaza, Thanks for the paper excerpts. I've already read it. The paper is just a belated admission of basic settled science on turbulent CFD simulations. Every honest scientist already knew this, except perhaps some climate scientists unfamiliar with science fundamentals. With incredibly coarse grid sizes, most scales are unresolved and the results are likely to be wrong even on the resolved scales. But details will be totally wrong. Despite decades of "circling the wagons" that's now becoming obvious for example concerning the pattern of SST warming. Models miss that and thus miss the global warming rate as a consequence.

Energy balance methods tell us what we need to know about temperature evolution. I do find it instructive that Palmer and Stevens must be so careful to avoid angering concensus enforcers by saying that GCM's may be badly flawed but global warming theory is secure albeit with broad ranges of uncertanty. That's all fine with me. The fact that you have lied here about my position on these fundamentals is very sad. You should retract your slander.

However, it is also settled science that even eddy resolving simulations might give poor results, It is settled science that classical methods of numerical error control fail for these simulations. In other words, its impossible to separate model error from numerical error and so tuning can't be completely scientific.

A computer programmer like yourself I know has a vested interest in the "Moon shot" billion dollar investment in eddy resolving GCM's. There is little real mathematics or science behind this. It's just vague cultural beliefs that "more physics" must be better. The problems with these simulations should be resolved (if its possible) using vastly simpler canonical problems.

I do think the evidence for this statement is weak.
"When it comes to global climate change, it is what the present generation of comprehensive climate models do not show—namely, a sensitivity of global changes to either the vagaries of unpredictable regional or global circulations or effects of processes neglected in simpler models—which makes them such a powerful confirmation of inferences from basic physics and simple models."
If details are wrong, GCM's do not show that those details are unimportant in the big picture. In fact its wrong with regard to patterns of SST warming as a score of papers have shown.

I do think that many climate modelers do not know the advantages of more recent numerical advances that can make simulations more accurate. Most GCM's still use finite volume methods and old time integration schemes. That's the situation in every field of CFD. Workhorse codes tend to use 30 year old numerical methods. That is a problem and I was correct about it. Snide insinuations is all you seem to have.

The leapfrog scheme used in many climate models is a vastly inferior and out of date method popular in the 1960's. Everything I have said is true. You have nothing in this regard as anyone reading this will see.

You really are guilty here of exaggerating what I've said and really misrepresenting it. It's sad.

Andrew on **May 11, 2020 10:06 PM at 10:06 pm** said:

David:

I recommend avoiding statements such as, "Every honest scientist already knew this." Lots of scientists are ignorant of all sorts of things. That doesn't mean they're dishonest.

I say this for two reasons. First, it's rude to call someone dishonest if they're not. Second, the equation of error with dishonesty has the problem that scientists who are not dishonest then assume they can't be in error. Unfortunately, life isn't so simple. A scientist can be honest and still be wrong.

dhogaza on **May 11, 2020 11:20 PM at 11:20 pm** said:

DK Young

"except perhaps some climate scientists unfamiliar with science fundamentals"

In addition to what Andrew said, perhaps you should stop suggesting that scientists who disagree with you are unfamiliar with science fundamentals?

"The paper is just a belated admission of basic settled science on turbulent CFD simulations"

If this were true, they wouldn't be suggesting a path forward being the use of massively more capable computers, far beyond supercomputer scale, in order to implement models with finer resolution so more CFD simulation can be done for various processes, and less parameterization.

The paper is an admission – which no one would agree with – that the currently available computer power to effectively increase the resolution of models to the point where parameterization can be abandoned for more processes that can be directly modeled through CFD simulations, or for processes that don't fit that model, by stochastic modeling, does not exist.

Thus the call for massive funds for building and making available to climate scientists massively more powerful computers.

Anyone can read the paper for themselves and see that it's not an admission of what you suggest are limitations on turbulent CFD. The paper suggests the the failure to more accurately model regional features are due to resolution being limited due to available computer resources.

"A computer programmer like yourself I know has a vested interest in the "Moon shot" billion dollar investment in eddy resolving GCM's."

"A computer programmer like yourself…" Ahhh the belittling labeling meant to bolster your argument from authority.

"There is little real mathematics or science behind this. It's just vague cultural beliefs that "more physics" must be better."

"I do think the evidence for this statement is weak.
"When it comes to global climate change, it is what the present generation of comprehensive climate models do not show—namely, a sensitivity of global changes to either the vagaries of unpredictable regional or global circulations or effects of processes neglected in simpler models—which makes them such a powerful confirmation of inferences from basic physics and simple models."
If details are wrong, GCM's do not show that those details are unimportant in the big picture."

So you praise Palmer and Stevens when it suits you, tear them down when you don't. You brought them into the discussion when you didn't anyone would actually read their paper.

Then, afterwards, you throw them under the bus.

Classic DK Young.

Hopefully this exchange might help people evaluate DK Young's objectivity regarding his efforts to convince people that the covid-19 epidemic is not of serious consequence, unless you're on your deathbed, and in that case, it doesn't matter …

David Young on May 12, 2020 12:46 AM at 12:46 am said:

I agree Andrew that ignorance does not imply dishonesty. One challenge in modern science particularly one as complex as climate science is that most are unfamiliar with broad parts of the fundamentals underlying their tools. That's not their fault either. Most CFD experts even with 40 years of experience are lucky to understand the fine details of turbulence modeling. The problem comes when a CFDer ignores a turbulence modeler when he says his model is very bad in massively separated flows and tries to insist that the model is indeed accurate by cherry picking results. That's a common problem in science generally. Better scientists will engage the experts and listen to what they say seriously. What DhoGaza is doing here is just trying to discredit that expert opinion without really contradicting anything the expert says.

I think DhoGaza that you have said nothing to really cast doubt on anything I said. Certainly there is nothing technical in your comment.

I don't think you read my comment carefully enough.

1. No one knows if massively increased resolution will result in skillful predictions. That's not even known for LES (which is what they are really proposing) for much simpler model problems. That's not well understood by many outside the turbulence modeling community. So its a common cultural meme to believe that if we "got all the physics in there" the results must be accurate. That's quite false.

2. What Stevens and Palmer are admitting is that parameterized models like current GCM's are quite bad at getting details correct, which is what all CFDers have known for 50 years. It's based on fundamental math and error propagation analysis.

These are two independent facts. Linking them as you do is a red herring.

You didn't really deal with any other technical issues I raised. Does that mean you agree that the leapfrog scheme is a very old scheme that has serious flaws such as nonlinear instability that are vastly improved by more modern methods like backward differencing methods? If you don't know its fine to say that.

Similarly finite element methods are much preferable to finite volume methods. There is a growing body of evidence showing that.

David Young on May 12, 2020 1:01 AM at 1:01 am said:

DhoGaza says:

"Hopefully this exchange might help people evaluate DK Young's objectivity regarding his efforts to convince people that the covid-19 epidemic is not of serious consequence, unless you're on your deathbed, and in that case, it doesn't matter …"

Misrepresenting my views and words is really dishonest and you should be ashamed. I never said any such thing. Covid19 is a serious epidemic. I believe however that's its not nearly as serious as people think and a lot of experts agree with this assessment. I agree that our best policy is to protect those who are vulnerable. These people do matter. I've been citing facts and data and expert opinion. You have been attacking my personally and lying about what I have said. It's a shameful display.

Carlos Ungil on May 11, 2020 5:51 PM at 5:51 pm said:

David,

In rough numbers, there are 55'000 deaths per year in NYC. There were 33'000 deaths in the last couple of months, 24'000 more than expected.

Let's say 2/3 of those excess deaths (i.e. 16'000) are from people who were going to die anyway within one year.

In the next ten months one would expect to see 30'000 deaths instead of 46'000 deaths (a 35% reduction).

Is that a fair representation of your position?

Reply ↓

David Young on May 11, 2020 6:13 PM at 6:13 pm said:

Yes, That would be a logical consequence of the estimate being correct. There might be confounding factors such as excess deaths from suicide or substance abuse given the high stress levels people are being exposed to.

dhogaza on May 11, 2020 11:23 PM at 11:23 pm said:

Carlos

Looks like we have a testable hypothesis. Though I don't expect that DK Young's going to care in 10 months, and very much doubt he'll be back here if he's wrong. Whatever policy decisions are made will mostly have played out in that timeframe, and policy is his real concern.

David Young on May 12, 2020 1:10 AM at 1:10 am said:

Mind reading again on display. A dishonest rhetorical tactic.

Everyone cares about policy to some extent, but I'm more concerned with people being scared to death and unwilling to go to the hospital when they have a heart attack or stroke. I'm concerned with the delayed mortality from skipped "nonessential" activities like medical screening tests. Delaying dental cleaning long enough results in real mortality. Isolating people in their homes and removing their income is very stressful and will result in a lot of deaths too.

Carlos Ungil on May 12, 2020 2:38 AM at 2:38 am said:

Ok, thanks. I think there will be some offsetting but I would be surprised is it was so high. By now, “they are dying with cover and not from cover” and “they would have died in a few weeks anyway” have been ruled out. Eventually we’ll find out if in most cases death was advanced by only a few months.

Zhou Fang on May 11, 2020 6:41 PM at 6:41 pm said:

Joshua:

I think that sort of motivated unapologetic incompetence becomes indistinguishable from malice.

David:

Do *you* have any substantive science? Because right now the entire basis of your argument is shown to be based on what might charitably be considered a case of misremembering. If you want to make this case why should other people dig through these papers for you?

Also yes if Ioannidis is claiming this that discredits him also.

Reply ↓

Zhou Fang on May 11, 2020 6:43 PM at 6:43 pm said:

(oops misreplied, ignore)

Joshua on May 11, 2020 7:04 PM at 7:04 pm said:

Zhou –

> I think that sort of motivated unapologetic incompetence becomes indistinguishable from malice

Sure. Practically speaking maybe you can’t tell and maybe there’s no real operational difference. Point taken.

David Young on May 12, 2020 9:37 AM at 9:37 am said:

Well Andrew, Zhou’s point is right that the 2/3 number is only an estimate. I also saw it elsewhere but can’t find it right now. This is not however the main point I made. That point is that the serious people think in terms of excess mortality,.

What surprises me is that you allow people to throw around ad hominems like “a pattern of deliberate deception.” This Is untrue and needs to be corrected.

I would like to find accurate numbers on nursing home fatalities but haven’t had time to do so. The Ioannidis paper on the Italian data is also worth a thorough read as I’ve seen numbers suggesting that the median age of fatalities is 81.

Reply ↓

Carlos Ungil on May 12, 2020 9:59 AM at 9:59 am said:

> The Ioannidis paper on the Italian data is also worth a thorough read as I've seen numbers suggesting that the median age of fatalities is 81.

I don't think anyone here disputes that it may be the case. That's also around the median age of all-cause mortality.

Is that supossed to support your estimate that 2/3 of those who died would have died anyway within twelve months? The life expectancy at 80 is closer to 10 years than to 1 year.

Carlos Ungil on May 10, 2020 5:16 AM at 5:16 am said:

> There are by now at least 10 serologic data sets around the world. They pretty much uniformly show an IFR less than 0.5% with the best ones showing perhaps 0.12% to 0.31%.

With 0.25% population fatality rate in NYC you need an infection rate over 50% for the IFR to be less than 0.5%. Does that look plausible to you?

This study was published this week: https://www.medrxiv.org/content/10.1101/2020.05.02.20088898v1.full.pdf

The population fatality rate in Geneva is 0.05% and they estimate prevalence to be 9.7% (+/- 2.5%). That would give an IFR of 0.54% (0.40%-0.85%).

But they also find that prevalance is 40% lower in the 50+ population, which is the relevant one regarding mortality (it could be even lower for the 65+ or 80+ but unfortunately we don't have better information). If we repeat the calculation with a more meaningful 40% lower prevalence we get an IFR of 0.9%. (The population is only slightly older than in the US, 16.5% over 65 vs 15.2%).

Reply ↓

dhogaza on May 10, 2020 4:46 PM at 4:46 pm said:

Carlos Ungil

But these aren't the "best" studies, as defined by David. The criteria for "best" I'll leave to your imagination.

An IFR of 0.5%-1% just doesn't fit the narrative.

Reply ↓

confused on May 11, 2020 8:30 PM at 8:30 pm said:

I agree the IFR is probably not below, say, 0.3%, but I'm not sure it would make sense to adjust upwards for the older population if the goal is to get a "nation overall" IFR.

It seems entirely expected that older people (outside LTCFs) would have a lower rate of infection, since many are retired and many are also likely to be more cautious (because it is more dangerous for them).

If this went to herd immunity, I wouldn't expect to see an equal percentage infected in all age groups.

If anything, this probably shows that an "overall" IFR is not that useful, or at least cannot be extrapolated from one region to another with different demographics, different activity profiles, etc. (In which case serological studies in New York or California don't tell us much about what the IFR in Texas or Alaska or South Dakota would be. I mean, it's probably not going to be 10x different, but I think we already know it to *that* degree of accuracy...)

Reply ↓

Carlos Ungil on May 12, 2020 8:36 AM at 8:36 am said:

I agree, the IFR stuff is complex. Assuming equal infection rates for all the population may not be realistic, but neither it is realistic to think that the distribution in the early stage is representative of the later evolution. And whatever the true IFR is now or later in NYC, it's not and will be not below 0.25%.

The slower the diffusion of the infection through the population the more homogenous it may be in the end. The original UK approach of massive infection while keeping the people at risk locked makes some sense (but there are too many uncertainties around it to say it was a good idea). By the way, older people are more cautious but maybe not as much as they should (many are complaining about discrimination).

Reply ↓

confused on May 12, 2020 5:46 PM at 5:46 pm said:

Sorry, I guess infection rate was the wrong word… I meant % of population infected (in that age group), not infections per day.

I was saying that a larger percentage of the younger population is infected, and a smaller percentage of the older population, the "population average" IFR would be lower than if % infected was equal for all age groups. And that it seems plausible that the "elderly but living independently" group would have a comparatively low % infected.

Reply ↓

confused on May 12, 2020 5:55 PM at 5:55 pm said:

Sorry, that comment was meant to be a response to Joshua's below.

Yeah the IFR in NYC will probably be close to 1%. I don't think it is going to be representative of other places in the US, though.

I don't think the US average, or even the "US excluding New York and New Jersey" average, IFR will be as low as 0.25%, no.

But I do think it will be noticeably lower than NYC's – because NYC and some other Northeast places seem to have done a particularly poor job at protecting LTCFs; because they were the first really severe outbreak in the US and their supportive care was probably worse than what's now practiced (better criteria on who to ventilate, proning patients, etc.); because while their hospital system wasn't actually overwhelmed to the "turning away patients" level, quality of care likely dropped due to overwork; and maybe, more speculatively, because more New Yorkers got a higher viral dose due to extreme population density, mass transit, etc.

I have wondered about stay-at-home for at-risk populations-only myself. I mean, if you don't have to go to work and can get groceries left outside your door, it should be totally possible to avoid any face-to-face contact with anyone who doesn't live in your house for the time span it would take for an unrestrained epidemic to go to herd immunity and die out.

But it couldn't be done just by age, it would have to also include anyone who lives with an elderly person as well (and same for immunocompromised only).

The question then would be – would an otherwise-unrestrained epidemic that infected, essentially, only young-to-middle-aged, non-immunocompromised people overwhelm hospitals?

Joshua on May 12, 2020 9:11 AM at 9:11 am said:

confused –

> It seems entirely expected that older people (outside LTCFs) would have a lower rate of infection, since many are retired and many are also likely to be more cautious (because it is more dangerous for them)

Unless I misunderstand something, the rate of infection is not really directly relevant to the IFR. although it might be indirectly so as a function of viral load.

Reply ↓

Joshua on May 12, 2020 9:14 AM at 9:14 am said:

I suppose it could also be indirectly relevant as a moderator in the relationship between treatment efficacy and mortality rate. IOW, a higher infection rate could lead to hospitals and healthcare workers being overwhelmed.

Reply ↓

confused on May 12, 2020 5:46 PM at 5:46 pm said:

Sorry, I guess infection rate was the wrong word… I meant % of population infected (in that age group), not infections per day.

I was saying that a larger percentage of the younger population is infected, and a smaller percentage of the older population, the "population average" IFR would be lower than if % infected was equal for all age groups. And that it seems plausible that the "elderly but living independently" group would have a comparatively low % infected.

Joshua on May 10, 2020 8:05 AM at 8:05 am said:

David –

> There are by now at least 10 serologic data sets around the world.

Respect uncertainty.

Santa Clara?
A hotspot.

Barcelona?
A hotspot.

NY?
A hotspot.

LA?
A hotspot.

Miami-Dade?
A hotspot.

Chelsea, MA?
A hotspot

Gangelt, Germany?
A hotspot

Hmmm. Maybe there's a pattern here?

Let's take the study in Chelsea, MA. You've mentioned that study before. You indicated it prives a high infection rate, a large denominator.

They estimated a 31.5% infection rate in Chelsea. Chelsea has 40,000 people. Right now, some 124 people from Chelsea have died from COVID (a number that's sure to go up). You do the math.

Here's the results of a recent meta-anslysis:

–snip–

> After exclusions, there were 13 estimates of IFR included in the final meta-analysis, from a wide range of countries, published between February and April 2020. The meta-analysis demonstrated a point-estimate of IFR of 0.75% (0.49-1.01%) with significant heterogeneity (p<0.001).

Respect the uncertainty.

Reply ↓

confused on May 10, 2020 4:25 PM at 4:25 pm said:

I agree with the general thought here, but I don't think Santa Clara or Los Angeles (or anywhere in CA) or even Miami-Dade County are really "hotspots".

Relative to some super rural and/or isolated (AK, AR, HI, UT, WY) part of the US, maybe, but in terms of deaths/population Miami and LA are not that bad compared to other very large urban areas in the US. The real hotspots among very large cities would be NYC/NJ (in a class of its own) and then Boston area (including Chelsea), Detroit, Chicago etc.

Chelsea, Gangelt, and NY are definitely hotspots though.

But yeah, I really can't see the IFR being as low as 0.12%-0.2%. It probably *will* be higher in New York than in most of the US, because New York seems to have done a particularly bad job of protecting nursing homes, and because New York had a lot of cases early on when less was known so probably more ICU cases died than would die today, but I can't see that making a factor-of-5 difference.

Iceland does show an IFR of about 0.55% with PCR testing alone, and there is no one left in ICU and only 18 active cases – so I do think New York City is to some degree a genuine outlier in IFR as well as prevalence. Iceland has the best PCR testing in the world, but I doubt they caught every single asymptomatic case. The Icelandic population may be a bit healthier than the US population, though.

Reply ↓

Carlos Ungil on May 10, 2020 5:02 PM at 5:02 pm said:

They have recorded only ten deaths in Iceland, so any estimate will come with huge error bars. It's remarkable (but likely to be just noise, don't get me wrong) that four out of twelve patients aged 80-90 died, while all the eight patients older than 90 survived.

Reply ↓

Joshua on May 10, 2020 7:00 PM at 7:00 pm said:

confused –

> I agree with the general thought here, but I don't think Santa Clara or Los Angeles (or anywhere in CA) or even Miami-Dade County are really "hot spots."

Miami-Dade? 14,000 out of 40,000 in the state?

https://www.google.com/amp/s/www.local10.com/news/local/2020/05/08/coronavirus-in-florida-latest-numbers-as-south-florida-reopening-mapped-out/%3foutputType=amp

The Silicon Valley in general and Santa Clara in particular had a relatively high concentration of infections early on – but like all of CA, the infection didn't take off after the shelter in place was initiated – similar to LA. Nonetheless, it would be a hotspot in terms of seroprevalence surveillance conducted when those studies were conducted.Certainly relative to other places even in CA let alone other parts of the country. Again, I'm talking about in terms of representativeness for extrapolating infection and fatality rates more broadly.

Same with LA. I'm not talking about hotspot in the sense of a lot of deaths over a longer period of time.

Reply ↓

confused on May 11, 2020 5:25 AM at 5:25 am said:

Miami-Dade is a hotspot relative to the rest of Florida, but it's not really a hotspot compared to other 2+ million population urban areas. Even given meatpacking plants, rural areas are generally doing much better, so 'largest city in state vs. whole state' comparisons are always going to show that the big city is worse.

Yeah, maybe the US average hadn't hit 2-3% seroprevalence or whatever at the time those samples in CA were taken. It might have, though, if there was really community spread in January…

I wonder what the US overall is now? 1.33 million confirmed cases is about 0.4% of the US population. If we're undertesting by 10x-20x as seems to be quoted a lot, that would suggest 4%-8% infected. But that 10x-20x is probably much too high now in many states (testing has improved a lot in some states over the last 2 weeks or so), and was probably much too low in February and early March outside possibly CA/WA…

Reply ↓

Joshua on May 11, 2020 8:56 AM at 8:56 am said:

confused –

> Miami-Dade is a hotspot relative to the rest of Florida, but it's not really a hotspot compared to other 2+ million population urban areas.

That seems a little silly to me. We may as well say that NYC isn't a hotspot relative to other cities with 18 million.

confused on May 11, 2020 8:48 PM at 8:48 pm said:

Well, New York is its own thing, not really comparable to any other US city. (Not just raw population size of the metro area, either – it's more mass transit/less car-friendly, incredibly dense by US standards, etc.)

But Miami-Dade County is 490 deaths out of 2.7 million, a bit over . That's worse than Houston or Dallas, but it's not really particularly high compared to the major cities in the Northeast and Midwest: Chicago, Boston, Detroit, Philadelphia. Even some smaller cities like New Orleans and Pittsburgh are doing far worse per-capita.

I think this is more a semantic argument than a substantive one, though. I'd agree that Miami is probably a hotspot relative to *the South* (though not nearly as bad as New Orleans), as most of the South except Louisiana has been fairly lightly impacted, but not relative to *the US as a whole*.

Joshua on May 12, 2020 9:03 AM at 9:03 am said:

confused –

I was being sarcastic with the NYC comment. Perhaps that didn't come across.

Once again, the death rates, particularly when considered over a longer period of time, is not the frame that applies to my use of the term "hotspot."

I was using that term w/r/t representativeness as a sample for extrapolating an infection rate, at the time the seroprevalence data were collected.

My point was w/r/t the context of the likely range of IFR (see my discussion with David), when estimated based on the use of seroprevalence data for extrapolation. In that sense, the Miami-Dade data would not likely be suitably representative for broad extrapolation, although it might not be that unrepresentarive for projecting rates for large cities in the US.

Zhou Fang on May 12, 2020 9:54 AM at 9:54 am said:

The thing is, one might argue that NYC is actually the most representative situation for our purposes because we are most interested in what the IFR would look like under easing of the lockdown and increase of attack rate.

confused on May 12, 2020 6:03 PM at 6:03 pm said:

@Zhou Fang… maybe. But I think NYC is *so* unlike the rest of the US that even with relaxed measures we wouldn't expect anything like NYC in most other places. NYC is just an extreme outlier among US cities in so many ways.

Some states never had stay-at-home orders, and mostly aren't doing clearly worse than other states (Iowa being a possible exception, depending on how fast its outbreak grows). The states that are relaxing measures aren't going to 100% pre-COVID conditions.

And most of the other cities I'd expect to be — while not comparable to NYC, nowhere is — at risk for the same sorts of reasons have already been hit pretty hard, so are likely to relax slower (and seroprevalence might be high enough to reduce effective R somewhat also).

confused on May 12, 2020 6:11 PM at 6:11 pm said:

Joshua –

Yeah, OK, I can see that. It's probably better not to draw sweeping conclusions from any particular seroprevalence study, and even with lots of them, to be careful about extrapolating to places that are very different in urbanization, demographics, etc.

IMO, we need seroprevalence studies in cities in the central US and in rural areas. I think there was one in Idaho, but…

I really wouldn't be surprised if 6% of the total US population had been infected, though. That would be 20 million people, and we have 1.36 million confirmed cases. Given just how bad testing was most places until the last ~2 weeks, real infections being 15x confirmed cases doesn't sound implausible.

On the other hand, 50-85x from the original Santa Clara study would mean 68-115 million people, or 20-35% of the US population. I don't think there's any chance the US overall seroprevalence is higher than NYC!

Zhou Fang on May 13, 2020 5:05 AM at 5:05 am said:

@confused:

But there's the thing, you have to go with a fixed notion of what "ending the lockdown" looks like. You could counter by claiming that even without the lockdown people will continue to stay at home but in that situation ending the lockdown will not have large economic benefits and might actually have a cost (due to lack of coordination and increased uncertainty).

confused on May 13, 2020 7:01 PM at 7:01 pm said:

> "But there's the thing, you have to go with a fixed notion of what "ending the lockdown" looks like."

The thing is that I don't think we will have anything like that.

Each state will do its own thing, except for those states that are explicitly coordinating on how to re-open, and even those will likely not do exactly the same thing.

In some places slight changes in the measures will be (are being, I think) hailed as "reopening" when they really don't change that much. Other places may reopen too fast (though it has probably been long enough for Georgia to blow up, and it doesn't seem to be doing that… I actually thought that one was higher-risk than, say, Texas.)

As for economic damage from fear vs. the formal measures – there will be for a while. After that it will depend on what actually happens. If we do not have a really bad outbreak (the definition of "really bad" probably being quite subject to political/media spin) in the next month or so, people will probably go back to something much closer to "normal".

Not really big mass events like sports with 50,000 fans in a stadium, but unless the summer is unexpectedly bad, I think by 4th of July people's restaurant-going behavior will be fairly normal, I think many (most?) schools and colleges will be open in the fall semester, etc.

Anoneuoid on May 13, 2020 7:36 PM at 7:36 pm said:

> *As for economic damage from fear vs. the formal measures – there will be for a while.*

I think more people are scared of being contact traced, quarantined, or medically detained in some way rather than of the virus.

So if that continues you can expect continued fear and further economic destruction.

Daniel Lakeland on May 13, 2020 8:18 PM at 8:18 pm said:

Confused: GA did their "open up" around April 23. Certainly people didn't just "go back to normal" so undoubtedly it would take 2 weeks minimum before we saw much change in movement, and it also takes typically 2 weeks from people start spreading it to show up in the daily numbers of cases. So, I really think we're going to see what GA teaches us *starting* next week and moving towards Jun 1. May 25-29 will be the week we start learn whether GA is going to be in trouble soon or not. By week of Jun 1 to Jun 5 we will be able to estimate the rate at which things are/aren't diverging. It will be informative.

Joshua on May 13, 2020 8:27 PM at 8:27 pm said:

Anoneuoid –

> I think more people are scared of being contact traced, quarantined, or medically detained in some way rather than of the virus.

What havd you seen that makes you think that's the case? How could you evaluate the ratio of people who feel one way or the other?

Joshua on May 13, 2020 8:37 PM at 8:37 pm said:

confused –

> I really wouldn't be surprised if 6% of the total US population had been infected, though

Given the range of estimates, inconsistency in test quality, lack of understanding of what levels of which antibodies means what w/r/t COVID-19 vs other coronaviruses, overall questionable representativeness of the sampling and variabliliry in the sampling, and the difficulty of accurately extrapolating from relatively small samples, I don't know how anyone can have any confidence about widely applicable infection levels except that they're somewhere between pretty low and pretty high.

IMO, it's one of those situations where we're probably better off saying "Yup, there's a lot of uncertainty."

Joshua on May 13, 2020 8:42 PM at 8:42 pm said:

Notice that the Stanford folks, experts who are deeply involved in these estimates, were apparently off in their estimates of how many MLB employees would test positive.

Anoneuoid on May 13, 2020 8:49 PM at 8:49 pm said:

> *What havd you seen that makes you think that's the case? How could you evaluate the ratio of people who feel one way or the other?*

Almost everyone I know, old and young, democrat or republican, think there is some kind of scheme behind this. I'd say 80%+, even 90%+.

confused on May 13, 2020 9:42 PM at 9:42 pm said:

@Daniel Lakeland: I guess, but I was seeing a lot of predictions that Georgia, Texas, etc. would blow up "in 2 weeks" (i.e. 2 weeks after opening). We're about 3 weeks in now.

Yes, I agree that Georgia isn't totally "out of the woods" by any means, but what I'm seeing now in their data is better than I expected, and what I expected was better than what a lot of people were predicting.

I am now thinking the worst is probably over in the vast majority of the US, at least for this "wave".

...assuming this pandemic even displays a wave pattern. But respiratory viruses do tend to be seasonal, so it would seem the "Occam's razor" assumption would be that this one will be also.

(I've seen it claimed that Brazil shows it's not seasonal, because Brazil is warm. But Brazil does have seasonal influenza, though it's not as strongly seasonal as the US, and it's now mid-late fall in Brazil, getting into their flu season. So I don't think that is an useful counterexample.)

Martha (Smith) on May 12, 2020 7:31 PM at 7:31 pm said:

This whole discussion reminds me of the old phrase, "The Map Is Not The Territory". Things like IFR are in some sense like "maps" — they are intended to describe some aspect of some situation the "territory") in some way. But just like maps, they may be constructed in a variety of ways, looking at the situation from different perspectives, including focusing on some small part of the the territory, or on a very large territory. I think someone said it already in the thread: "Remember the uncertainty" (or, probably better: "Remember the variability". One size (e.g, one IFR) does not fit all "territories".) So it makes sense to consider a range of plausible IFR's for whatever "territory" is being described.

Reply ↓

confused on May 13, 2020 9:46 PM at 9:46 pm said:

@Joshua: Yes, there's definitely a lot of uncertainty. I wasn't saying that I'm confident the real level is around 6%; the only results that would actually surprise me would be 20% (in which case the whole US has been as heavily infected as NYC, which would require NYC to have a ridiculously higher IFR – I think somewhat higher is pretty likely, but not *that* much).

Reply ↓

confused on May 13, 2020 9:48 PM at 9:48 pm said:

Ugh, messed up that comment. I meant that the only results that would really surprise me would be 20% (given that's about where NYC seems to be).

Joshua on May 13, 2020 10:21 PM at 10:21 pm said:

confused –

Regarding the level of uncertainty:

> One study released Wednesday examining an Abbott Laboratories test that's used at the White House to get rapid diagnoses indicated it may miss as many as half of positive cases. A second peer-reviewed study released hours later suggested that results for another type of widely used diagnostic test are particularly unreliable early on in an infection.

https://www.msn.com/en-us/money/markets/false-negatives-raise-more-questions-about-virus-test-accuracy/ar-BB143341

confused on May 13, 2020 11:48 PM at 11:48 pm said:

And somehow my comment still didn't post correctly… I believe the site mistook less-than and greater-than signs for HTML tags.

What I meant to say was that I would only be surprised by the true seroprevalence being *less than* 2% (which would mean we'd caught more than 1/5 of all infections) or *greater than* 20% (comparable to NYC).

The test being criticized there is a PCR test, not serology. And false negatives would mean we are probably missing even more cases than we thought, in which case only 2%-3% or so of the US population being infected (when confirmed cases are more than 0.4%) seems even less plausible.

But yeah, tons of uncertainty.

Joshua on May 14, 2020 12:01 AM at 12:01 am said:

confused –

Of course I realize it isn't the seriology test – but the problem with the false negatives in the PCR isn't the only implication w/r/t uncertainty. It also adds to the problems with determining infection fatality rate.

I didnt post that to suggest that the uncertainty only runs in the direction of reducing the infection rate – but to intricate that people who are making certain claims either way, that restrict probability rages at this point, are likely displaying "motivated" reasoning.

confused on May 14, 2020 5:04 PM at 5:04 pm said:

Yeah, OK, I can see that.

I do think the issues with testing in the US in March and April, combined with the number and wide distribution of the cases we *did* find, do kind of set a lower bound on how low of a prevalence is plausible, though.

Joshua on May 14, 2020 7:11 PM at 7:11 pm said:

confused –

Was it with you thst I was discussing CDC estimate of flu IFR?

–snip–

The root of such incorrect comparisons may be a knowledge gap regarding how seasonal influenza and COVID-19 data are publicly reported. The CDC, like many similar disease control agencies around the world, presents seasonal influenza morbidity and mortality not as raw counts but as calculated estimates based on submitted International Classification of Diseases codes.2 Between 2013-2014 and 2018-2019, the reported yearly estimated influenza deaths ranged from 23 000 to 61 000.3 Over that same time period, however, the number of counted influenza deaths was between 3448 and 15 620 yearly.4 On average, the CDC estimates of deaths attributed to influenza were nearly 6 times greater than its reported counted numbers. Conversely, COVID-19 fatalities are at present being counted and reported directly, not estimated. As a result, the more valid comparison would be to compare weekly counts of COVID-19 deaths to weekly counts of seasonal influenza deaths.

During the week ending April 21, 2020, 15 455 COVID-19 counted deaths were reported in the US.5 The reported number of counted deaths from the previous week, ending April 14, was 14 478. By contrast, according to the CDC, counted deaths during the peak week of the influenza seasons from 2013-2014 to 2019-2020 ranged from 351 (2015-2016, week 11 of 2016) to 1626 (2017-2018, week 3 of 2018).6 The mean number of counted deaths during the peak week of influenza seasons from 2013-2020 was 752.4 (95% CI, 558.8-946.1).7 These statistics on counted deaths suggest that the number of COVID-19 deaths for the week ending April 21 was 9.5-fold to 44.1-fold greater than the peak week of counted influenza deaths during the past 7 influenza seasons in the US, with a 20.5-fold mean increase (95% CI, 16.3-27.7).5,6

https://jamanetwork.com/journals/jamainternalmedicine/fullarticle/2766121?guestAccessKey=28f9727c-0a86-4d26-8900-e4391ec18af2&utm_source=silverchair&utm_medium=email&utm_campaign=article_alert-jamainternalmedicine&utm_content=olf&utm_term=0051420

Phil on May 17, 2020 2:39 AM at 2:39 am said:

Anonueioud says

> I think more people are scared of being contact traced, quarantined, or medically detained in some way rather than of the virus.

and

> Almost everyone I know, old and young, democrat or republican, think there is some kind of scheme behind this. I'd say 80%+, even 90%+.

I've just gotta mention the Pauline Kael fake quote, "I can't believe Nixon won. I don't know anyone who voted for him." She did not in fact say this, what she said was "I live in a rather special world. I only know one person who voted for Nixon": she was well aware that the people she knew were not a representative sample of society.

Anon, I have friends who are wary of contact tracing but none who have expressed outright opposition, and none who oppose quarantine for infected people. You and I know very different groups of people. I know my friends are not a representative cross-section of America. Do you know that yours aren't either?

Reply ↓

Ben Bolker on May 10, 2020 6:35 PM at 6:35 pm said:

I'm not sure which comment of Dalton's you're responding to. Excess mortality (and lives-lost as raised in Phil's new post on this blog) is certainly worth thinking about when considering impact (not sure where 333,000 -> 13% excess mortality comes from? Is 295K the number of expected deaths?). On the other hand, if you're comparing predicted vs observed for COVID19 mortality I don't see why you'd want to use anything but mortality …

For what it's worth, it seems that excess mortality analyses are mostly pointing in the other direction, i.e. that COVID19 mortality impacts are underestimated by COVID death counts. There are obviously a bunch of tricky issues here (e.g. (1) "harvesting" – as you suggest, how many of the deaths are accelerating deaths that would have occurred soon anyway? (2) how do we count a heart attack death that occurred because someone didn't go to the hospital out of fear of COVID19? Is that "caused" by fear, lockdown, COVID19 … ?)

Reply ↓

David Young on May 11, 2020 10:54 AM at 10:54 am said:

If you use the Miami Dade numbers, the excess mortality if 300 million are infected is 6% – 11% of annual mortality which is roughly 2.8 million. It's not a small number, but its also not the end of the world. Looking at case fatality rates is just so despicably wrong and alarmist that it makes me angry that the media do this. Case fatality rates range from 2% to over 10% depending on which state you look at.

Reply ↓

Carlos Ungil on May 11, 2020 11:17 AM at 11:17 am said:

Where are these numbers published? What is the estimated prevalence by age?

For the hypotehtical case where the whole US population is infected, why would these numbers be more applicable that those of NYC where 0.25% of the population has died? (33k people have died in the last two months, when less than 10k deaths would have been expected.)

Reply ↓

dhogaza on May 11, 2020 11:53 AM at 11:53 am said:

DKY sez:

" Looking at case fatality rates is just so despicably wrong and alarmist that it makes me angry that the media do this"

It's all they've had to report on. They've also reported on statements by Ioannidis and his colleagues at Stanford that "it is no worse than the flu", and that the IFR might be as low as 0.12%.

Reply ↓

David Young on May 11, 2020 4:16 PM at 4:16 pm said:

Real journalists would report real estimates of infections and deaths, not the phony case numbers. They really are being irresponsible and frightening poeple.

I've seen a couple of interviews with Ioannidis but they were brief compared with the wall to wall coverage of misleading numbers, anecdotal stories or horrible scenes of refrigerator trucks outside hospitals or mass graves. It's been the worst yellow journalism episode since the 19th century.

Reply ↓

David Young on May 11, 2020 11:17 AM at 11:17 am said:

For the uncertainty mongers, there is also data from Arizona. It appears that the calculated IFR using current deaths is 0.21%. As in my Miami Dade calculations, one should use a fatality number from 2 weeks after the mean testing date. If we assume a 30% increase in deaths over the next week or two as happened in Miami Dade, that gives 0.28% right in the middle of the range from Los Angeles, Santa Clara, and Miami Dade.

Its obvious with a moment's thought (except perhaps to computer programmers) that IFR will vary a lot between localities because the percentage of the vulnerable who are infected will vary a lot. In some places, more than 50% of fatalities have been among nursing home residents. In those places IFR's will be higher. Florida seems to have done a good job with nursing homes. New York did a disastrous job and so their IFR will be higher.

I would argue that New York is an outlier because of the very poor response of state and city authorities with regard to protecting vulnerable people. Most of the country in terms of population density, use of mass transit, etc. is closer to Miami Dade, Los Angeles, and Santa Clara, and Arizona. Most of the better run states have probably done a better job with nursing homes too.

It is also obvious that many many young people have had this disease and had a very mild course. This is why Josh your silly IFR "symptomatic" is such a silly number that no one else uses but the authors of your silly study. There are millions of people who had a very mild course. Doctors would disagree about whether they were "asymptomatic" or not. It's just a number that is meaningless in a practical sense. You keep conflating a true IFR with this IFRAsynpotomatic pseudo statistic. It just confuses people.

Reply ↓

David Young on May 11, 2020 11:44 AM at 11:44 am said:

Sorry, should have said IFR Symptomatic. In any case, its an impossible to accurately measure quantity.

Reply ↓

dhogaza on May 11, 2020 11:57 AM at 11:57 am said:

DK Young sez:

"Its obvious with a moment's thought (except perhaps to computer programmers) that IFR will vary a lot between localities because the percentage of the vulnerable who are infected will vary a lot"

I actually think we agree that the optimistic IFR numbers coming out of Stanford shouldn't be extrapolated to the country at large. Combined with uncertainty, basing policy on these numbers would be stupid. And the PR push to do so was irresponsible.

Reply ↓

> David Young on May 11, 2020 4:22 PM at 4:22 pm said:
>
> As I said, Miami Dade, Los Angeles, Arizona, and Santa Clara are much more representative of the US as a whole than New York City or Chelsea Mass.
>
> What PR push? It's completely responsible to push for re-opening the economy give the massive costs of keeping it shut down. Delaying medical screenings and treatments (as has happened across the country) will result in lots of future fatalities. Poverty and unemployment kill too.
>
> Reply ↓

David Young on May 13, 2020 12:55 PM at 12:55 pm said:

I'll just add another observation here. What happened to John Fund and then to me here shows the toxic nature of anonymous internet commenters who have no real knowledge of the subjects being discussed but have strong political emotions. They inevitably resort to the most gross form of ad hominem insults and lies about what their targets have said. I think Andrew, requiring people to use their real names might improve this childish behavior. It's just like school yard bullies except that these bullies can weak a hood.

On the science, you will note that Dhog and Joshua have provided nothing. Dhog regurgitated the contents of a paper that I referenced. The paper said exactly what I said it did. He has provided absolutely nothing else but slanders and lies. I'm surprised that you allow it.

Reply ↓

> Anoneuoid on May 13, 2020 1:33 PM at 1:33 pm said:
>
> Welcome to the internet.
>
> Btw, after seeing some of your other posts... do you have a comment on this: https://statmodeling.stat.columbia.edu/2020/05/10/uncertainty-and-variation-as-distinct-concepts/#comment-1334183
>
> No one has spent the time to correct me.
>
> Reply ↓
>
> > dhogaza on May 13, 2020 1:51 PM at 1:51 pm said:
> >
> > If you want to be corrected, why not run your ideas by practitioners in the field? David Young has been unable to get traction with his, shall we say, somewhat unique ideas and stoops to misrepresenting papers by others in an effort to bamboozle people on the internet (welcome to the internet, indeed). I trust that you're more honest and if practicing climate scientists point out to you why you are incorrect about GCMs, for instance, that you'll take it as a learning opportunity rather than insist you must be right, and them wrong.
> >
> > Right? You'd do that, right?
> >
> > Reply ↓
> >
> > > Anoneuoid on May 13, 2020 2:08 PM at 2:08 pm said:
> > >
> > > Sure, point it out to them. Anyone is free to respond.
> >
> > Joshua on May 13, 2020 2:15 PM at 2:15 pm said:

Anoneuoid –

>No one has spent the time to correct me.

Since you're interested in a feedback, you might be able to get some at this site:

https://andthentheresphysics.wordpress.-com/

Or this one:

http://julesandjames.blogspot.com/

dhogaza on May 13, 2020 2:16 PM at 2:16 pm said:

James Annan, who works on climate sensitivity (narrowing uncertainty, in particular), has posted here a few times in regard to covid-19.

You could try contacting him and run your ideas past him. Though he might not be willing to engage …

Anoneuoid on May 13, 2020 2:20 PM at 2:20 pm said:

Im not going to be bothering anyone. They can respond there if they see it.

Joshua on May 13, 2020 2:26 PM at 2:26 pm said:

Anoneuoid –

> Im not going to be bothering anyone. They can respond there if they see it.

Joshua on May 13, 2020 2:30 PM at 2:30 pm said:

Anoneuoid –

> Im not going to be bothering anyone. They can respond there if they see it.

I'm reasonably sure that practically no one over at this site:

https://andthentheresphysics.wordpress.-com/

Would consider it a bother. There are a lot of people who frequent that site who like to talk about climate modeling. The person who runs the side specifically does so to talk about climate science with people who are interested in exchanging views in good faith (although sometimes, no doubt, people part ways in terms of establishing what good faith is).

They aren't mostly climate modelers themselves, but fairly knowledgeable about such and sometimes climate modelers do read the comments section. Give it a shot and see what happens. Can't hurt.

Joshua on May 13, 2020 2:31 PM at 2:31 pm said:

I'll tell you what. If you'd like I'll post it for you – for "a friend" – and see if anyone bites… So then I'll be the person running the risk of bothering someone. Just let me know if that would work for you.

Anoneuoid on May 13, 2020 2:51 PM at 2:51 pm said:

@Joshua go for it. You can link to that post or just repost it without attribution, doesnt matter much to me.

David Young on May 13, 2020 2:56 PM at 2:56 pm said:

Dhog, Over the last 20 years, there is a growing numbers of papers showing that GCM's have the same limitations as other areas of application of CFD. If you really have something

David Young on May 13, 2020 2:59 PM at 2:59 pm said:

Dhog, Except that there is a growing corpus of papers showing that climate models suffer the same limitations as other areas of application of CFD. Stevens and Palmer is just the latest. I mentioned some in my response below. If you have something that has actual technical substance, I'd entertain it. But please drop the slander.

Joshua on May 13, 2020 3:04 PM at 3:04 pm said:

Anoneuoid –

Will do. I"ll let you know if there are any bites:

BTW – from the paper that David said he represented accurately:

> We know that greenhouse gases are accumulating in the atmosphere as a result of human activity and that they are largely responsible for warming of surface temperatures globally. We also are confident in our understanding as to why this warming is expected to be amplified over land masses and the Arctic. Likewise, we are confident in our understanding of how the hydrological cycle amplifies the effects of this warming and how warming amplifies the hydrological cycle. For these and other broad brush strokes of the climate change picture, we are also increasingly confident in our ability to usefully bound the magnitude of the effects. From this certainty stems the conviction that additional warming is best avoided by reducing or reversing emissions of long-lived greenhouse gases.

As climate scientists, we are rightfully proud of, and eager to talk about, our contribution to settling important and long-standing scientific questions of great societal relevance. What we find more difficult to talk about is our deep dissatisfaction with the ability of our models to inform society about the pace of warming, how this warming plays out regionally, and what it implies for the likelihood of surprises. In our view, the political situation, whereby some influential people and institutions misrepresent doubt about anything to insinuate doubt about everything, certainly contributes to a reluctance to be too openly critical of our models. Unfortunately, circling the wagons leads to false impressions about the source of our confidence and about our ability to meet the scientific challenges posed by a world that we know is warming globally.

——

Not it's rather interesting that he mentioned the "circling the wagons" part but neglected to mention the rest. Very curious, that.

David Young on May 13, 2020 3:21 PM at 3:21 pm said:

Josh, What you cite is not new or interesting. What I emphasized is vastly more important and consequential because it has consequences for future research.

Joshua on May 13, 2020 3:22 PM at 3:22 pm said:

Anoneuoid:

Done:

https://andthentheresphysics.wordpress.com/2020/05/09/attacking-scientists-who-tell-the-truth/#comment-175993

Joshua on May 13, 2020 3:26 PM at 3:26 pm said:

David –

> What you cite is not new or interesting. What I emphasized is vastly more important and consequential because it has consequences for future research.

Ok. Got it.

You selectively left all that off because it's vastly less important and interesting. And of course, you are the arbiter of importance ant interest and so no one else should find that other part interesting or important either. So you were actually just performing a service for them by editing out the uninteresting and unimportant parts.

The fact that those other parts don't align well with your frequent critique of the work of climate scientists, I'm sure, has absolutely nothing to do with the mechanics of your editing.

Got it.

Joshua on **May 13, 2020 3:36 PM at 3:36 pm** said:

Andrew –

If you happen to read this, I apologize for cluttering the "recent comments" list with this childishness. Bad form on my part. Your blog deserves better.

I'll stop now. (Although I will find Anoneuoid and post any responses to his comment if they come through).

Anoneuoid on **May 13, 2020 3:52 PM at 3:52 pm** said:

> *Anoneuoid:*
>
> *Done:*
>
> *https://andthentheresphysics.wordpress.com/2020/05/09/attacking-scientists-who-tell-the-truth/#comment-175993*

Thanks, I will try to keep an eye on it. I don't think it was understood so far from the response though. I am not making some theoretical argument. It is simply logical deduction from how they have defined it.

David Young on **May 13, 2020 5:05 PM at 5:05 pm** said:

Josh anyone who wants can read the paper in full. It's hard in a blog comment to present everything. I don't know what the purpose of you quoting a few sentences is. My characterization of the paper is correct and anyone can check for themselves.

Anoneuoid on **May 14, 2020 3:06 PM at 3:06 pm** said:

> *Anoneuoid:*
>
> *Done:*
>
> *https://andthentheresphysics.wordpress.com/2020/05/09/attacking-scientists-who-tell-the-truth/#comment-175993*

There was more discussion but not a single person addressed the point. They just go on thinking the range represents an uncertainty...

David Young on May 13, 2020 2:54 PM at 2:54 pm said:

Anon, Your comment is perceptive in my opinion. Climate sensitivity is not a system constant and does vary perhaps a lot depending on the current state of the system.

That said, I doubt that a lot of progress on narrowing the IPCC range of values is likely. The reason for this is that there is too much reliance on GCM's (General Circulation Models) in keeping the upper part of the range alive.

The bottom line fact here is that GCM's are weather models run with large grid sizes (about 50km) and running them for centuries.
1. On these grid sizes, most important scales are not resolved and so "subgrid models" of things like turbulence, tropical convection, clouds, etc. are required.
2. These subgrid processes are themselves ill posed in an important sense and the subgrid models are well knows to be very inaccurate (even turbulence modeling which is the oldest one are well known to have serious deficiencies) and to require vastly finer grids than those used in the GCM's.
3. The ECS involves small changes in energy flows to be accurately modeled, such as clouds. These small changes are orders of magnitude less than the numerical errors.
4. Skill is achieved on some output measures such as global mean temperature anomoly by tuning the parameters of the code (and there are hundreds of them). Basically, top of atmosphere radiation balance is tuned and ocean heat uptake is roughly right. Since energy is conserved, average temperature change will be ball park right.
5. Other outputs such as ECS will be all over the place. There are some recent papers showing that changing some details in the convection and cloud models within the data uncertainty (which is large) result in quite significant changes in the ECS of the model.
6. Yet another paper is more disturbing. It shows that changing the order of application of the subgrid models results in large changes in output too.
7. Despite these issues, its not even clear if super fine grids would result in adequate skill. The proposal of Palmer and Stevens is really Large Eddy Simulation. This method is just now starting to be evaluated by the CFD community and its unknown at the moment if it will prove skillful. We know that grid convergence is questionable. More seriously classical methods of numerical error control fail. Thus it is impossible to distinguish numerical error from errors in the other errors in the model itself. There may be ways around these issues, but it requires fundamental theoretical research.

I cannot emphasize too strongly that every one in CFD has known about these issues with resolution for 60 years going all the way back to Von Neuman. No one would try to do an engineering calculation with these very coarse grids and try to assert that the calculation meant much in the context of engineering design for example.

There are huge and very interesting pure research issues here to be resolved here. The best theoretical people should be working on them and get funding priority. Just building ever more complex models will accomplish virtually nothing other than making Dell and Intel more profitable.

Further, by giving the impression (circling the wagons) that these are "solved" problems, the funding stream for further theoretical progress has dried up. This is a really bad result of the general and well documented strong bias in science to generate positive results and file negative ones in the recycling bin. "After all, my bad results were due to lack of my skill in choosing grid sizes, setting parameters, etc." The supply of witches is endless. These excuses are honestly believed but biased.

It is a vast problem that I think is causing progress in this area to stall. Honesty and a very strong attempt to overcome biases (including appropriate training of students) is needed. That's why I really admire John Ioannidis and trust him. He has been in the fore front of exposing the strong biases in medicine and trying to point the way forward. He is just fantastic, humble, and always open to opposing points of view. It is very discouraging to see Gavin Schmidt for example smearing him on twitter.

Reply ↓

dhogaza on May 13, 2020 3:29 PM at 3:29 pm said:

I thought you were going to actually link to some papers so we can check to see if you're accurately representing their findings, as you've done with Palmer and Stevens, and to see if your claims that GCMs are doomed to failure are supported by these papers.

On a gross level, nearly everything your list is known and acknowledged by the climate modeling community, so it doesn't support your position that they're ignorant of them or intentionally ignore them.

Anoneuoid on May 13, 2020 3:50 PM at 3:50 pm said:

> *Anon, Your comment is perceptive in my opinion. Climate sensitivity is not a system constant and does vary perhaps a lot depending on the current state of the system.*

It isn't just that though. It is that the way they define it in the equations it is a variation (like the standard deviation), **not** an uncertainty (like a confidence interval). So the whole goal of reducing it is misguided (and impossible) to begin with.

David Young on May 13, 2020 4:12 PM at 4:12 pm said:

Dhog, You are mis-stating what I said yet again. What I do believe is exactly what Palmer and Stevens point to. I believe Schmidt knows many of these facts. He is probably not an expert on turbulence models and their strong limitations. What happens is that to avoid acknowledging error or uncertainties, the community (which is strongly in favor of action) resorts to all kinds of gymnastics to avoid admitting publicly what they may privately acknowledge. Schmidt is unusual in that he has strong mathematical training. Many of those who just run the models are not familiar with the underlying facts of CFD. It should be people like Schmidt who insist on the limitations and advocate that people need to do better. Circling the wagons is a terrible, terrible strategy on every level.

This happens in most fields of science. Turbulence modelers all know that their models have very serious deficiencies in separated flows or even in strongly 3D flows. Yet most practitioners of CFD don't fully understand these limitations. Modelers don't want to risk disrupting the funding stream by highlighting the deficiencies. But this strategy has actually backfired very badly.

Funders are now convinced that CFD is a solved problem. Funding for fundamental research has dried up and the result is disastrous for science. In my view, this is a deep cultural problem within science itself and will require fundamental reforms. One culprit is the growing reliance on soft money even at top research universities. In the old days tenured professors had their research hours paid by the institution and they also could get funding for graduate students or even postdocs. Now, they have to get grants to survive. The temptation to oversell is very strong.

David Young on May 13, 2020 4:15 PM at 4:15 pm said:

Yes Anon, I agree. What is clearly bad statistics is to treat models as if they represent a random sample of possible models and then compute a mean and a standard deviation. This method almost certainly underestimates uncertainty, but its the best we can do without tremendous effort. I've done it myself in CFD but did say that it underestimated uncertainty.

David Young on May 13, 2020 4:18 PM at 4:18 pm said:

Dhog, I have never said GCM's are doomed to failure. What is true is that progress will be slow and painful. New strategies will be needed and new ways of communicating the needs for fundamental research. Please stop paraphrasing in single simple-minded and inaccurate phrases.

Anoneuoid on May 13, 2020 4:24 PM at 4:24 pm said:

> *Yes Anon, I agree. What is clearly bad statistics is to treat models as if they represent a random sample of possible models and then compute a mean and a standard deviation. This method almost certainly underestimates uncertainty, but its the best we can do without tremendous effort. I've done it myself in CFD but did say that it underestimated uncertainty.*

Eh, apparently no one understands what I am saying. The range given for the climate sensitivity (1.5-4.5C) is like a standard deviation, not a confidence interval. It is a descriptive property of the data so the goal of making that range smaller makes no sense. Spending money for that purpose is literally throwing it down a logical black hole.

Atomsk's Sanakan (@AtomsksSanakan) on June 21, 2020 2:49 PM at 2:49 pm said:

David, you've been rebutted on your GCM point for awhile now, and you have no cogent response (no, refusing to refuse the published literature is not a cogent response):

http://www.realclimate.org/index.php/archives/2019/12/how-good-have-climate-models-been-at-truly-predicting-the-future/comment-page-2/#comment-752062

> "That's why I really admire John Ioannidis and trust him. He has been in the fore front of exposing the strong biases in medicine and trying to point the way forward. He is just fantastic, humble, and always open to opposing points of view. It is very discouraging to see Gavin Schmidt for example smearing him on twitter."

Your hero worship is noted. As I've told you before on Curry's blog, Ioannidis does not agree with you on climate science. He's stated that the level of certainty on humans causing climate change is about on par with the idea that smoking is killing people, and that climate science contains reproducible results that justifies decisive government action on climate change:

"Many fields lack the high reproducibility standards that are already used in fields such as air pollution and climate change. [...]
It is a scandal that the response of governments to climate change and pollution has not been more decisive."
https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5933781/

17:17 to 18:30 :
http://rationallyspeakingpodcast.org/show/rs-174-john-ioannidis-on-what-happened-to-evidence-based-med.html
http://hwcdn.libsyn.com/p/f/8/9/f89584694d653aa6/rs174.mp3?c_id=13577443&cs_id=13577443&expiration=1584648769&hwt=736ba1cc30e54af900f17b584b4685ec

In any event, David, do you have a cogent reply to GidMK's threads rebutting Ioannidis' under-estimate of IFR? Here are GidMK's threads on this, along with two other threads discussing studies he hasn't included in his meta-analysis yet:

on version 2 of Ioannidis' review:
https://threadreaderapp.com/thread/1270490491600003072.html
on version 1 of Ioannidis' review:https://threadreaderapp.com/thread/1262956011872280577.html
https://twitter.com/GidMK/status/1267683223712104450 (with: "A systematic review and meta-analysis of published research data on COVID-19 infection-fatality rates")

https://twitter.com/GidMK/status/1267669313491632129
https://twitter.com/AtomsksSanakan/status/1272722887967944709

If that's too broad a topic for you, then explain why Ioannidis took a Brazilian study with a stated IFR of 1.0%, and decreased that by a factor of ~4 to ~0.27%. I see no good reason for him to do that, and it just looks like motivated reasoning (at best) on his part:

"Median 0.30 (0.27)"
https://www.medrxiv.org/content/10.1101/2020.05.13.20101253v2.full.pdf

"The ratio of deaths over estimated cases – the estimated infection-fatality rate – was 1.0% (7,640)."
https://www.medrxiv.org/content/medrxiv/early/2020/05/30/2020.05.30.20117531.full.pdf somebody on May 13, 2020 2:36 PM at 2:36 pm said:

What happened to John Fund is that he said a bunch of demonstrably false and misleading statements, then people were rude when they called him on it.

1. Do you dispute that reporting the upper point of a confidence interval as a prediction is misleading
2. Do you dispute that reporting projections up to 2080 as someone's failed prediction is misleading
3. Do you dispute that comparing someone's projections under a counterfactual which didn't take place to the observed numbers is misleading?
4. Do you dispute that multiplying a case fatality rate by the total infection count to infer someone's deaths is misleading?

I don't think you do. I'm not really sure how people being rude to him shows their lack of real knowledge of the subjects being discussed. It seems like your problem there at least is exclusively with the tone.

On your own comments:

1. You also reported the upper bound of someone's estimates as their projection, and it wasn't even a formal upper bound but rather an offhand comment about how little we know. Charitably, this is you misremembering what you saw somewhere.
2. You've asserted with strangely high confidence that a very large percentage of excess mortality would have died within the year without evidence other than that there's a strong correlation between age and fatality. You cite that the median age of the virus's victims are 80, but the conditional life expectancy of someone at 80 is 10 years, not one.
3. When people pointed out your assertions about IFR don't make sense for NYC, you said that won't be representative for the rest of the country because of vague claims like the policy there was "bad." I'm actually rather curious what specifically you're referring to when you say NYC's policy on nursing homes was bad, or that its governance has been incompetent.
4. You claim that Josh's meta-analytic review is of a "symptomatic IFR", which as far as I can tell is not true. I looked at the review and as far as I can tell these are estimates of true IFR, and studies reviewed include a serological study and a bunch of modeling based estimates of the true population prevalence including asymptomatic. Are you just lying here?
5. You ignore another serologic dataset that doesn't conform to your claim.

On top of that, you editorialize your reporting of a range of serologically derived IFRs by saying that "the best" studies show the lowest IFRs, without explaining at all how you're judging quality.

My point is not that you're wrong. Maybe we should pull a Sweden, maybe IFR in most of the U.S. will be 0.1 percent. But it really looks like you're misrepresenting your level of certainty with vagueness and misrepresentation. People pointing these things out aren't nitpicking; the details are relevant. Maybe people are ruder than they should be, but the fact remains that you're making claims that are stronger than your evidence. So I would expect some pushback.

Reply ↓

David Young on May 13, 2020 3:19 PM at 3:19 pm said:

I don't think I misrepresenting anything. My memory is not perfect with regard to the 2/3 number. It's was just an offhand estimate. As I said, I did hear it somewhere else, I think in an Ioannidis video. He has several long ones from the last 2 months. I'm trying to do some research to establish a better number. If you want to help rather

than just nipping at my heals, I'd like that.

I do think there is a lot of uncertainty here and what I've cited tries to estimate these. The main point here is just that the IFR numbers are vastly lower than the very alarming and completely meaningless the CFR's parroted by the media.

I also think Santa Clara, Los Angeles, Arizona, and Miami Dade are more like the rest of the US than New York City. In addition New York did a terrible of protecting nursing homes and that may have contributed to a skewed IFR. But we need better data to know for sure. If you want to help look into that I would like that.

If you have something to add that adds to my understanding that would be great. I'm not going to respond further to your microscopic but vague assertions about what I said above. People can read it for themselves if they want to check.

Reply ↓

somebody on May 13, 2020 4:06 PM at 4:06 pm said:

> If you want to help rather than just nipping at my heals, I'd like that.

I'll be forthright in saying that I don't have anything to add on the substantive science. But other people have had something to add on the substantive science and you've dismissed them just because you don't like their tone or because you perceive them to be part of an angry mob against you.

1. Someone pointed out that the life expectancy at 80 is 10 years, which suggests that perhaps you're overestimating the share of excess mortality which will happen within the year
2. Someone linked to another serological study that estimates IFR at above 0.5%
3. Joshua linked to a meta-analytic review of IFR studies, your criticism of which I already pointed out is demonstrably false. Not to say that it's a good review–I think the whole enterprise of trying to compute "the" infection fatality rate like it's planck's constant is misguided. But it's points at a lot of papers with different approaches to estimating total population prevalence.

I do think I can help you to understand why you're being so negatively received: it's that you don't get to make strong or suggestive claims and then "do some research" to find out where they come from afterwards. You also don't get to tone police people who call you out for not having a source or an argument, then condescendingly make fun of people while simultaneously being wrong. You say my criticisms are "microscopic yet vague." So I'll be more specific for you.

In your first comment

You say

> 2. There are by now at least 10 serologic data sets around the world. They pretty much uniformly show an IFR less than 0.5% with the best ones showing perhaps 0.12% to 0.31%.

This is the kind of vaguely suggestive editorializing that sets of alarm bells for me. It seems like you're trying to say without saying that the IFR is correctly bounded at 0.3%. In any case, you don't get to make a claim about what the best studies are saying without explaining why they are the best. Furthermore, there are more than 10 serological datasets, many of which show an IFR greater than 0.5%. Another way of saying what you said is that "I picked 10 serological data sets which show an IFR less than 0.5% and the ones I like the most show 0.12% to 0.31%". Which is honestly fine to say, but it's easy to misread your rhetorical presentation as being essentially "the most reasonable boundary on IFR is under 0.31%."

I'll take a moment to say that I agree that the CFRs being reported on are probably misleading the public. Even if I agree with your substantive point though, I think your rhetoric here is irresponsible.

> 3. Total fatality numbers are misleading and would not be used by competent scientists. The number that is useful is the total excess mortality number. By Ferguson's own estimation, 2/3 of those who die with covid19 would have died within the year because they were already seriously ill.

At the very least, you saying Ferguson's estimate is 2/3rds is misleading. At this point, you maybe, maybe don't have another source for the number, but you've already made a claim about what someone said which is not true. In response to people calling you out about the claim, you said they were nitpicking.

You also said this

> It is also obvious that many many young people have had this disease and had a very mild course. This is why Josh your silly IFR "symptomatic" is such a silly number that no one else uses but the authors of your silly study. There are millions of people who had a very mild course. Doctors would disagree about whether they were "asymptomatic" or not. It's just a number that is meaningless in a practical sense. You keep conflating a true IFR with this IFRAsynpotomatic pseudo statistic. It just confuses people.

Which was both incredibly condescending and completely wrong. AT BEST you got confused about your open tabs and at worse you lied and hoped nobody would actually check.

Again, my point is that people are justified in being negative towards or skeptical of your claims. Even if your main point that CFRs reported by the media are misleading or alarmist is correct, and I think it is, you don't get to be sloppy about the argument, then mad when people point out that it's sloppy. People pointing out that it's sloppy aren't demonstrating "a lack of real knowledge of the subjects being discussed but strong political emotions". Being careful about epistemology and methodology is the entire point here.

> In addition New York did a terrible of protecting nursing homes and that may have contributed to a skewed IFR.

I have a genuine question here that you seem to have missed. What, specifically, do you think New York did poorly which other states will be able to handle better without a lockdown? I'm in New York right now and policy towards those with comorbidities seems pretty tight.

somebody on May 13, 2020 4:18 PM at 4:18 pm said:

> The main point here is just that the IFR numbers are vastly lower than the very alarming and completely meaningless the CFR's parroted by the media.

Also want to point out that I don't think anyone would really dispute that IFRs are more meaningful than CFRs and reporting CFRs is misleading the public. All the disputes have been with other information you added in to buttress the point which didn't really check out. Throwing in extra weak arguments with a strong one doesn't make the whole statement stronger, it makes it weaker. The weak arguments don't not matter so long as your point is broadly correct.

dhogaza on **May 13, 2020 4:20 PM at 4:20 pm** said:

Preliminary results of Spain's national serological testing study, consistent with an IFR of about 1.1%. Yes, Spain's health care system was overwhelmed (saving David the trouble of pointing that out).

https://www.lamoncloa.gob.es/servicios-deprensa/notasprensa/sanidad14/Paginas/2020/130520-seroprevalencia.aspx

Regarding New York's "terrible of protecting nursing homes and that may have contributed to a skewed IFR", their reported figure ALL long-term care facilities (nursing homes are a subset) is 20% of the state's total deaths.

As you point out, David said: "By Ferguson's own estimation, 2/3 of those who die with covid19 would have died within the year because they were already seriously ill".

And now say "My memory is not perfect with regard to the 2/3 number. It's was just an offhand estimate"

David put a lot of blog posting capital into defending that 2/3 number, as can be seen by anyone reading his comments. Maybe everything he so strenuously defends is really just an "offhand estimate"?

Carlos Ungil on **May 13, 2020 5:02 PM at 5:02 pm** said:

> Preliminary results of Spain's national serological testing study, consistent with an IFR of about 1.1%.

Even higher when we take into account the undercounting. In Madrid, the most affected region, 11.3% (9.8-13.0) prevalence means 750k people infected. Confirmed deaths are 8.760, almost all of them at hospitals, implying a IFR of 1.2% (1.01-1.34). The official estimate from the regional government including those who died at retirement homes is over 15'000, rasing the IFR to 2% (1.7-2.3). Maybe the prevalence in the older population is higher (for the whole country they give estimates by age and for all the age groups older than 40 the point estimate of the prevalence is above the population mean) but in any case the IFR seems very likely to be over 1%.

http://cadenaser00.epimg.net/descargables/2020/05/13/749ec6d73a8a14c1ed389711079cbfe5.pdf

David Young on May 13, 2020 9:19 PM at 9:19 pm said:

No one should be surprised that different localities have vastly different IFR's. Even using Ferguson's IFR's which are probably substantially too high.

Under 40 — IFR ~0.029%
Under 50 — IFR ~ 0.066%
Over 60. — IFR ~ 5%

If everyone in the US under 40 got infected, there would be perhaps 29,000 who would die according to these estimates. If everyone over 60 was infected, there might be 5,000,000 deaths which is a huge number. That's why its a scandal that states didn't do a better job of protecting nursing homes. In a lot of places over 50% of fatalities took place in these settings.

Does anyone know if the Spanish data was corrected to be age representative? Was their sample random?
I think the Santa Clara data was corrected. Miami Dade at least made an effort to do random testing. Any differential in the age distribution of those exposed will make a significant difference.

David Young on May 13, 2020 9:42 PM at 9:42 pm said:

Another school yard tactic to discredit and sow doubt from an anonymous commenter who cherry picks single sentences:

"Maybe everything he so strenuously defends is really just an "offhand estimate"?"

It's a transparently fact free insinuation with no merit. 2/3 is an estimate from Ferguson that I also saw in another source perhaps a month ago. I've looked at hundreds of articles on this over the last month and can't remember every single word of all of them. I would be interested in any real input on other estimates though.

As I said, this illustrates perfectly the nasty and shameful nature of so many anonymous commenters on the internet.

Joshua on May 13, 2020 10:40 PM at 10:40 pm said:

somebody –

BTW, re that "Cuomo killed all those old people in NY meme that's circulating among rightwing idealogues…

> As of Wednesday, Belgium, with a population of over 11.4 million, has counted a total 6,262 deaths from COVID-19 — roughly 540 per million citizens — and more than half of those deaths were in nursing homes. Of those 52%, just 4.5% were confirmed as having been infected, yet all are counted in the national tally.

Cuomo just be busy what with running NY and Belgium.

https://www.npr.org/sections/coronavirus-live-updates/2020/04/22/841005901/why-belgiums-death-rate-is-so-high-it-counts-lots-of-suspected-covid-19-cases

David Young on May 13, 2020 10:52 PM at

10:52 pm said:

Somebody, I'm not going to reread your list of very minor issues. In the time you wasted to type it you could have done some research and made a real contribution.

David Young on May 13, 2020 11:07 PM at 11:07 pm said:

I just realized that my estimate for those over 60 in the US is wrong. I forgot to take into account that there are many many fewer in each age category as age increases. Recalculating based on Ferguson's numbers, an accurate estimate would be 1.528 million, still a large number, but much much smaller than 5 million. Sorry for the miscalculation.

somebody on May 13, 2020 11:17 PM at 11:17 pm said:

David Young

2 out of the very first 3 things you said were completely wrong, and you either lied or misremembered when citing 2 sources on this page, you pee pee head. Making a large number of false claims, the falsehood of which you don't even dispute, isn't a "minor issue" just because your main point is obviously true, and then very publicly not caring about being demonstrably wrong over and over just ruins your credibility.

somebody on May 13, 2020 11:23 PM at 11:23 pm said:

And I'm still wondering what NY did wrong wrt to senior citizens! I have senior relatives in NYC right now, what is Cuomo doing to them? Should I know something?

Carlos Ungil on May 14, 2020 2:29 AM at 2:29 am said:

> And I'm still wondering what NY did wrong wrt to senior citizens!

I have not followed the issue, but this seemed a bad idea:

[March 26] New York Mandates Nursing Homes Take Covid-19 Patients Discharged From Hospitals

https://www.wsj.com/articles/new-york-mandates-nursing-homes-take-covid-19-patients-discharged-from-hospitals-11585228215

David Young on May 14, 2020 2:47 PM at 2:47 pm said:

Somebody, I don't care what an anonymous internet nobody with no qualifications says and I doubt anyone else cares. Getting a life would perhaps give you something more productive to do with your time.

Carlos Ungil on May 15, 2020 2:51 PM at 2:51 pm said:

> I have not followed the issue, but this seemed a bad idea:
> [March 26] New York Mandates Nursing Homes Take Covid-19 Patients Discharged From Hospitals

A recent article on the subject:
https://www.wsj.com/articles/new-york-sent-recovering-coronavirus-patients-to-nursing-homes-it-was-a-fatal-error-11589470773

New York Sent Recovering Coronavirus Patients to Nursing Homes: 'It Was a Fatal Error'

The state reversed its policy after mounting criticism and deaths. The mandate is part of broader scrutiny of weaknesses at long-term care facilities that have made them hot spots for Covid-19.

Atomsk's Sanakan (@AtomsksSanakan) on June 21, 2020 1:23 PM at 1:23 pm said:

> "No one should be surprised that different localities have vastly different IFR's. Even using Ferguson's IFR's which are probably substantially too high."

So David Young repeats the same ridiculous tactics here that he uses to misrepresent climate science (and medical science) over on Judith Curry's denialist blog. The difference here is that he doesn't have Judith Curry present to block or edit comments that rebut his claims, and he's dealing with a larger number of competent people who aren't falling for his tactics:

https://judithcurry.com/2020/05/06/covid-discussion-thread-vi/#comment-916987

Anyway, on to a few points.

It looks like Ferguson's model is reproducible:
https://www.nature.com/articles/d41586-020-01685-y

Ferguson's model used an IFR of 0.9%:
"These estimates were corrected for non-uniform attack rates by age and when applied to the GB population result in an IFR of 0.9% with 4.4% of infections hospitalised (Table 1)."
https://web.archive.org/web/20200609200001/https://spiral.imperial.ac.uk:8443/bitstream/10044/1/77482/14/2020-03-16-COVID19-Report-9.pdf

That's with the range of 0.8% – 1.0% IFR observed for the UK:

https://www.thelancet.com/pdfs/journals/lanpub/PIIS2468-2667(20)30135-3.pdf
https://twitter.com/GidMK/status/1268003875287490561 (with: "A systematic review and meta-analysis of published research data on COVID-19 infection-fatality rates")
https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7289569/

And that observed 0.8% – 1.0% IFR is with interventions that reduce IFR, such as lockdowns. So if anything, Ferguson more likely under-estimated the relevant IFR as opposed to over-estimating it like David claimed. Thus David was wrong, though I don't expect him to admit that, given his track-record. I think there was also a previous comment on here regarding Ferguson's model under-estimating COVID-19 deaths:

https://statmodeling.stat.columbia.edu/2020/05/08/so-the-real-scandal-is-why-did-anyone-ever-listen-to-this-guy/#comment-1332269

> "Does anyone know if the Spanish data was corrected to be age representative? Was their sample random?"

David, you can drop the illusion that consilient seroprevalence results support the false narrative you, Ioannidis, etc. have been peddling for months on low IFR. The above cited work from GidMK shows that, as does other material:

https://www.mdpi.com/2079-7737/9/6/128
https://threadreaderapp.com/thread/1270490491600003072.html

David Young on May 13, 2020 3:20 PM at 3:20 pm said:

I may be confused on which study Josh cited here. He has cited another one at Climate Etc.

Reply ↓

Joshua on May 13, 2020 10:04 PM at 10:04 pm said:

somebody –

I just want to say, I hope you know that you'll get nowhere in terms of trainable dialog.

Intersting study out of Denmark.

As a point of interest, compare the under 60 fatality numbers to those our friend has posted.

https://www.medrxiv.org/content/10.1101/2020.04.24.20075291v1.full.pdf

Reply ↓

Joshua on May 13, 2020 10:15 PM at 10:15 pm said:

Ha. That was supposed to be "reasonable" dialog.

Re that study in Denmark, I thought this was interesting:

> Conversely, blood donor prevalence increases with income and we speculate that this leads to higher risk of exposure through travel and social activity. We may therefore either under or overestimate the true population immunity.

———

IMO, probably doesn't work that way in the States. Sure maybe more travel with wealthy but poorer people more likely to have more exposure (use public transportation, be an essential worker) here.

David Young on May 13, 2020 10:46 PM at 10:46 pm said:

> "In this survey of SARS-CoV-2 antibodies in Danish blood donors we found a sero-prevalence of 1.7
> (CI: 0.9-2.3) adjusted for the assay performance and a low IFR of 82/100,000 (CI: 59-154). This
> IFR of 0.082% is slightly lower than a recently published COVID-19 IFR estimate of 0.145% (CI:
> 0.088-0.317, individuals below 60 years) not including seroprevalence data"
>
> That looks to me a lot lower than Ferguson's estimate for people under 70. Did you see any age distribution data Josh? This is actually the lowest credible estimate I've seen so far. But as I mentioned, IFR's will vary over a wide range depending on the age profile of those infected.

Joshua on May 13, 2020 3:14 PM at 3:14 pm said:

David –

> They inevitably resort to the most gross form of ad hominem insults and lies about what their targets have said. I think Andrew, requiring people to use their real names might improve this childish behavior. It's just like school yard bullies except that these bullies can weak a hood.

I will point up that frequently in your responses to me on any variety of issues (even ones which are not technical in nature) you counter my arguments with reference either to the fact that I post anonymously or that I'm not a scientist. In fact, we can see some examples in this very thread, and I would be happy to link many more examples where you make such statements if you think I'm wrong in my characterization.

Now one definition of ad hominem is when someone points to some aspect of the person making an argument in response to their argument, rather than responding directly to the argument. It's rather interesting that not only do you frequently employ that structure of argument directly in exchanges with me, but that you also (1) complain about ad homs and (2) do so under your full name while arguing that people posting anonymously is associated with the use of ad homs.

What makes it even more ironic in this case was Andrew calling you out for personalizing the discussions.

I would be more than happy to stick to non-ad hom exchanges with you in the future. I believe that I have mostly avoided ad hominem arguments when exchanging with you in the past (i.e., pointing out that there is a motivated bias in your reasoning is not that sort of ad hom although I can see where it borders on such; I will also note that usually I point out that you are not unique in that regard and that most people fall into that same trap) whereas you regularly employ ad homs in your exchanges with me. So my belief is that it would require less change on my part than it would for you.

But whatever, I will be more than willing to make sure I don't employ such a rheotorical frame in the future when exchanging views with you. I hope that you do so as well. Maybe the fact that you post under your full name will actually ensure that you're better at it than I am?

Reply ↓

David Young on May 13, 2020 3:59 PM at 3:59 pm said:

OK Josh. Thanks.

Reply ↓

somebody on May 14, 2020 2:50 PM at 2:50 pm said:

"I'll just add another observation here. What happened to John Fund and then to me here shows the toxic nature of anonymous internet commenters who have no real knowledge of the subjects being discussed but have strong political emotions. They inevitably resort to the most gross form of ad hominem insults and lies about what their targets have said. I think Andrew, requiring people to use their real names might improve this childish behavior. It's just like school yard bullies except that these bullies can weak a hood."

Reply ↓

Martha (Smith) on May 13, 2020 11:00 PM at 11:00 pm said:

David Young used "obvious" at least twice in his 11:17 am post. As a mathematician, I shudder –the word would get points off if a student used it in homework or on an exam. It's at best a sloppy or subjective term — also often used in the meaning of, "I am right, so if you disagree with me, then you must be wrong." I don't recommend it if you want to be credible.

Reply ↓

Atomsk's Sanakan (@AtomsksSanakan) on June 23, 2020 2:12 AM at 2:12 am said:

> "There are by now at least 10 serologic data sets around the world. They pretty much uniformly show an IFR less than 0.5% with the best ones showing perhaps 0.12% to 0.31%."
"For the uncertainty mongers, there is also data from Arizona. It appears that the calculated IFR using current deaths is 0.21%. As in my Miami Dade calculations, one should use a fatality number from 2 weeks after the mean testing date. If we assume a 30% increase in deaths over the next week or two as happened in Miami Dade, that gives 0.28% right in the middle of the range from Los Angeles, Santa Clara, and Miami Dade.
[...]
New York did a disastrous job and so their IFR will be higher.
[...]
Most of the country in terms of population density, use of mass transit, etc. is closer to Miami Dade, Los Angeles, and Santa Clara, and Arizona."

I've dealt with you on other forums before, David, so I know you often misrepresent science to suit your ideology. You're low-balling IFR to avoid lockdown policies you dislike, just like you low-ball climate sensitivity to avoid climate policies you dislike. In your discussions of IFR, you've committed or peddled a number of mistakes, including:

1) Under-estimating the rate of false positives in serologic testing.
2) Using a sample size or proportion of people who tested positive that's small enough to exacerbate the relative impact of under-estimating the test's false positive rate.
3) Using seroprevalence research on blood donors who are more likely to be healthy, skewing the estimated IFR lower than in a general population survey.
4) Using seroprevalence research on recruited people; they are more eager to enroll in the study because they're at greater risk of infection, inflating seroprevalence and thus skewing the estimated IFR lower.
5) Under-estimating COVID-19 deaths in comparison to excess mortality, skewing the estimated IFR lower.
6) Falsely claiming New York is extremely high outlier for IFR in the USA [discussed more below].

I've pointed out some of these issues to you before, and folks like GidMK covered them:
https://judithcurry.com/2020/05/06/covid-discussion-thread-vi/#comment-917143
https://threadreaderapp.com/thread/1262956011872280577.html
https://threadreaderapp.com/thread/1270490491600003072.html
https://judithcurry.com/2020/05/06/covid-discussion-thread-vi/#comment-917143
https://reason.com/2020/05/18/do-the-divergent-results-of-covid-19-antibody-studies-reflect-real-differences/

So I'll wrap up with an overview some USA IFR estimates, since you focused on the USA. And I'll include non-peer-reviewed results as well, since you cited non-peer-reviewed material, as per my above link to Judith Curry's denialist site. This overview should help illustrate how you under-estimated IFR.

When I say "GidMK review", I'm referring to at least one of these 3 sources/threads:
https://twitter.com/GidMK/status/1268003875287490561
"A systematic review and meta-analysis of published research data on COVID-19 infection-fatality rates"
https://twitter.com/GidMK/status/1267669313491632129

When I say "Ioannidis review", I'm referring to:
https://web.archive.org/web/20200611090545/https://www.medrxiv.org/content/10.1101/2020.05.13.20101253v2.full.pdf

Reply ↓

Atomsk's Sanakan (@AtomsksSanakan) on **June 23, 2020 2:16 AM at 2:16 am** said:

On to the list:

Serologic, New York:
1.4% (New York City) : https://www.worldometers.info/coronavirus/coronavirus-death-rate/
0.9% (New York City) : "SARS-CoV-2, COVID-19, infection fatality rate (IFR) implied by the serology, antibody, testing in New York City"
0.9% (New York City) : GidMK review
0.8% (state of New York) : GidMK review

Serologic, non-New-York:
0.8% (Chelsea, Massachusetts) : http://archive.is/UDrho [April 14/15: ~32% infected ; April 28: ~98 deaths ; http://archive.is/ymQ9khttps://web.archive.org/web/20200429145023/https://www.chelseaama.gov/coronavirusupdates]
0.6% (Indiana) : https://www.coronavirus.in.gov/2393.htm [https://news.iu.edu/stories/2020/05/iupui/releases/13-preliminary-findings-impact-covid-19-indiana-coronavirus.html]
0.6% (Arizona) : https://www.azdhs.gov/preparedness/epidemiology-disease-control/infectious-disease-epidemiology/covid-19/dashboards/index.php [June 7: ~3.1% infected ; June 21: ~1342 deaths ; http://archive.is/ymQ9k]

Reply ↓

Atomsk's Sanakan (@AtomsksSanakan) on **June 24, 2020 11:17 AM at 11:17 am** said:

Criticized and/or non-representative serologic results:
0.1% (Mission District of San Francisco, California) : "SARS-CoV-2 community transmission during shelter-in-place in San Francisco"

Non-serologic:
(values given for each USA state) : "Global prediction of unreported SARS-CoV2 infection from observed COVID-19 cases", figure 4 and supplemental table S1 [illustrates that New York was not an extreme outlier]
1.3% (symptomatic cases, USA) : "Estimating the infection fatality rate among symptomatic COVID-19 cases in the United States"
0.07% – 1.4% (USA) : "Using influenza surveillance networks to estimate state-specific case detection rates and forecast SARS-CoV-2 spread in the United States"
0.8% (North America [USA + Canada]) : "Predicted COVID-19 fatality rates based on age, sex, comorbidities, and health system capacity"

Reply ↓

More Anonymous on **May 8, 2020 2:36 PM at 2:36 pm** said:
I'm not sure who critics think would have done a much better job than Neil Ferguson.

Sure, there are probably some people who would have a done somewhat better job — it is after all improbable that the person who ends up with the most responsibility will be exactly the best persons available. But if there is someone out there who would do much better than Ferguson, I've certainly never heard of them.

I'm familiar with some of Neil Ferguson's papers due to research from years before Covid-19. Papers by him and co-authors (such as Christophe Fraser, Simon Cauchemez, and Christl Donnelly) are some of the absolute best work on rapidly developing epidemics that I found. When Covid-19 started and Neil Ferguson's name showed up in the news, I was delighted to hear people were listening to him. Its one of the few good things that have happened with Covid-19.

The work by him and the Imperial College of London teams continues to be the best modeling I've seen for Covid-19, overall. I've criticized aspects of it on this blog and disagree with parts strongly. But that's not the point. The point is there is no other group doing better work overall.

Reply ↓

SDE on May 8, 2020 2:48 PM at 2:48 pm said:

If his sims are the best, then . . .
The math in the papers may be fine, I haven't looked at it in detail. But I trust no output from this code base.

Who knows, his forecasts may be right, even his advice, but running that code doesn't tell us anything.

Reply ↓

More Anonymous on May 8, 2020 3:19 PM at 3:19 pm said:

There are many models from the Imperial College of London teams. I'm not sure if you dislike all of them, most of them, or one of them. If you dislike one of them, then can I suggest that you consider the value of them overall, rather than that one in isolation. If you dislike most of them, I'm not sure where you are coming from.

I expect there to be problems with the work from any group that goes from being unknown to at the center of catastrophe planning in a few weeks, especially when we are in a situation where speed must be prioritized because even a delay of a day or two in getting research out can have bad repercussions (more deaths). But to me the overall quality of work from the Imperial College of London teams is excellent.

Reply ↓

SDE on May 8, 2020 6:40 PM at 6:40 pm said:

My issue is with the code base that was open sourced, including its git history (which ends the discussion about MS somehow being the source of problems). I read a good chunk of it, and as this is what I do for a living (ML at a major tech firm) I have a pretty good idea of what makes code that generates reasonably trustworthy output. What I saw when I looked at the code base horrified me, as well as a wide swath of my engineering and science colleagues from around the company. I'm not worried about any one bug, and I don't care about some other random blog post. I'm worried about code that is so poorly written as to preclude any trust in the output. And if code of that quality is just normal research code in epidemiology, then the field has serious problems. Even when I was an academic, I would have been pretty bothered by this code, though not as much as I should have been. I appreciate that they were almost certainly doing their best, but I don't think the code had any value for forecasting. We were better off relying on exponential growth curves, plots of raw data, and back of the envelope estimates.

Reply ↓

More Anonymous on May 8, 2020 11:03 PM at 11:03 pm said:

I'm not really sure how to respond to you because you haven't responded to my earlier points:

First, the ICL team Ferguson leads has used many models (including with code on github) so if you are questioning the overall worth of the ICL team, then please consider their overall work rather than only attacking the one model you dislike the coding of.

Second, the authors were likely rushed because they rapidly went from being a relatively little-known team to central players in a situation where even short delays in reporting results would probably result in more deaths.

Yes, the code you dislike looks it could have been written years and years ago by people whose CS educations ended before Python and R existed. Maybe it was.

I get that the experience of seeing the code may have been very frustrating, in the same way as a doctor going back to see how medicine was practiced long ago would be horrified ("This is barbaric Jim, barbaric!"). But your criticisms just do not make me rethink the value of the ICL team's work.

Reply ↓

GCB on May 9, 2020 11:46 AM at 11:46 am said:

You are giving SDE too much credit here. Points made by dhogaza, somebody, and Zhou Fang throughout the comments (mostly in reply to Bob) have shown that both SDE and Bob are pushing falsehoods. And when pressed, all SDE can do is ridiculously turn up the "I'm an authority" knob. This post of SDE's is just more of the same.

Tangentially, I would like to point out that much (all?) of R is built on C, C++, and even Fortran; with a good deal of legacy code predating Python, etc. And Stan's core and math libraries are written in C++. The choice of language and even coding conventions should not make you rethink the value of its reliability.

Reply ↓

More Anonymous on May 9, 2020 5:40 PM at 5:40 pm said:

GCB: Thanks for your points. Yeah, you're probably right. I think it's best to assume good faith, even when one suspects moderate bad faith — but beyond a certain point, it's not worth it anymore.

I also should have been clearer about old programming styles and conventions — I agree the important thing is whether they work to purpose, not whether they look pleasing to the modern eye.

I don't think I implied that C, C++, or FORTRAN are bad, but if so I shouldn't have.

dhogaza on May 9, 2020 9:24 PM at 9:24 pm said:

GCB

Well, as far as good faith goes, this is so clearly right out of the play book of the climate science denialists who pounded for GCMs to be made public, and once it was done, in the case of NASA GISS's Model E, made the same bogus argument.

"Not written to industry standards"
"No professional software engineer would find this code acceptable"
and of course "it's written in FORTRAN!"

So I dug into the Model E sources and found them reasonably easy to read despite all the brou-ha-ha, despite not having been fluent in the language since FORTRAN 66, long before they dropped the all caps and called it Fortran.

Same shit. Recognized it right away.

dhogaza on May 9, 2020 9:28 PM at 9:28 pm said:

Interestingly, Bob and SDE have disappeared. Both are probably busy writing world-class software ...

GCB on May 10, 2020 11:43 AM at 11:43 am said:

More: Yeah, I should have said, “should not make *a person* rethink…”. Was really meant more as a general statement than criticism of you. I got that you weren’t saying the languages are bad. You are absolutely right about assuming good faith before bad. Something I am trying to work on!

dhogaza: Spot-on about the climate denial playbook BS. Appreciate your calling it out so effectively. Don’t fret, Bob and SDE will be back. If not in this thread then another.

David Young on May 11, 2020 1:39 PM at 1:39 pm said:

What constitutes “good” code is very subjective. Programmers all have their biases. This is a form over substance issue.

The problem here is the same as with GCM’s. The problem is ill-posed and in the case of gcm’s it’s well known that the simulations have huge levels of numerical and sub grid model9ng errors. Stevens and Palmer have a new paper basically acknowledging what every cfd er has known for 50 years. They get global average temperature more or less right due to tuning TOA fluxes and ocean heat uptake. But the patterns are wrong. For 30 years people like Hansen and dhog have engaged in a consensus enforcement campaign that Stevens andPalmer call “circling the wagons.” That’s a charitable way to put it,

When the problem is Ill-posed and millions of times more complex than you can model you will have quite large uncertainty. Better to invest in fundamental research into things like attractor dimension and Lyopanov exponents. Spending billions on more and more codes absent progress on fundamentals is a waste.

Reply ↓

yyw on May 8, 2020 3:04 PM at 3:04 pm said:

The difficulty in evaluating model performance based on predictions for inherently difficult to predict phenomenon like a pandemic is that you can make all kinds of errors in math and coding and these errors can make the results off by 100% or even an order of magnitude, but you would still be in the ballpark given that the ballpark is by necessity (unless someone shows otherwise) so huge.

Reply ↓

Joshua on May 8, 2020 4:13 PM at 4:13 pm said:

yyw –

> The difficulty in evaluating model performance based on predictions for inherently difficult to predict phenomenon like a pandemic is that you can make all kinds of errors in math and coding and these errors can make the results off by 100% or even an order of magnitude, but you would still be in the ballpark given that the ballpark is by necessity (unless someone shows otherwise) so huge.

What is your goal in your evaluation? If it is to determine whether the model is “correct,” then I’d argue that you’re chasing your tail.

Seems to me that the goal of evaluating a model should be to test the parameters that determine the range of uncertainty. If those parameters are correct, then the wide range is correct. A model with a wide range is neither incorrect nor not useful. It is what it is. If the parameters are uncertain, then they are what they are, and you have a wide range in your output You can’t correct for a wide range in a model’s output be injecting some kind of phony certainty into the parameters.

So if you want to evaluate a model, don’t look at the range. Look at the parameters. Correct them if you can. My assumption is that someone like Ferguson has a pretty solid understanding of the parameters.

Reply ↓

yyw on May 8, 2020 11:53 PM at 11:53 pm said:

You assume there is a consensus on a parametric modeling approach and that the only uncertainty is the values of the parameters. Model should always be evaluated by its output. You look at parameters to spot obvious errors.

Reply ↓

Ron Kenett on May 8, 2020 4:05 PM at 4:05 pm said:

I read the blog and some of the comments. There is one important thing that seems to have been out of the radar. In a tweeter exchange with Fergusson the issue of replicability came up. In response to a question he responded that his code was 13 years old and not documented so that, in fact, was not available for a replicability exercise. A critical point decision makers should be able to understand is that anyone providing advice should be subjected to a test of replicability. Fergusson did not pass this test…

Reply ↓

Zhou Fang on May 8, 2020 4:14 PM at 4:14 pm said:

This is why there's there's a coordinated research program to help document and re-implement the model?

https://epcced.github.io/ramp/

Reply ↓

Daniel Lakeland on May 8, 2020 4:05 PM at 4:05 pm said:

The bullshit asymmetry principle rides again!

Reply ↓

More Anonymous on May 8, 2020 4:22 PM at 4:22 pm said:

+1, also see "Gish Gallop"

And btw, thanks for your tireless persistence in the blog comments today.

Reply ↓

somebody on May 8, 2020 4:23 PM at 4:23 pm said:

Well, this is a shitshow. For the angry people in here to bash on Neil Fergusson who may or may not still be reading though, some reflection:

It seems like you've developed a narrative that the people here are blinded by some kind of ideological bias and are determined to smear your criticisms on the basis of your perceived agenda. The fact of the matter is that at least some of the criticisms that have been advanced here and in the original post are pretty unambiguously terrible. Some examples:

1. The original article at a couple of points argues that Fergusson wildly overprojected deaths by comparing his projections assuming no lockdown with real numbers under a lockdown.

2. The original article claims an overprediction from BSE by citing the top end of an interval of projections up to 2080.

3. People in the comments of the blog have repeatedly criticized Ferguson's code on the basis of a github repo which was not his code. In fairness, many of the same people have also acknowledged that it's not his code, but proceeded under the reasoning of "well if this is the later cleaned up version and it's still this bad, think about how bad the original must have been!" Maybe his original code was bad, but that logic is totally unsound. First of all, new bugs get introduced in the translation process from a one-time script to a product, the github repo is explicitly stated to be unfinished, and maybe the new software developers are actually worse than him? Again, I'm not saying his code isn't bad — it probably is, like most academic code — but rather that your criticisms which reference specific bugs, specifically attributed to Fergusson, is pointless.

4. People keep pointing out this one non-determinism bug which is almost certainly attributable to new features introduced in productization and pointing it back to Fergusson, chiding people for not reading the issue tracker, then contradicting information that's available in the issue tracker. No interest in fixing it? It's been fixed!

My point here is not to say that Fergusson's predictions are good. I have pretty fundamental issues with the whole framework of SIR and agent based simulation modeling, and personally prefer to go with a Talebian style reasoning of assuming a worst case on purely theoretical grounds when quality data for model falsification is limited. I also think that refusing to give out code on something as high stakes as this is pretty sketch, and that code quality standards in academia are pretty bad.

But if there are valid criticisms here, they're clearly being swamped with nonsense, and the critics seem to be buttressing each other at every turn anyways. It gives the impression that you have collectively developed an attitude of "it doesn't matter if I spit out 5 or 10 or 15 bad arguments so long as they're all pushing in the right direction. It doesn't matter if some of my logic is faulty or if I lied on the stand as long as I know that he's guilty." At that point, I'm not going to listen to you because who knows if you're saying anything worthwhile. If I snoop through on your first 3 sources and they're bogus but your 7 next sources are great and support your point, I'll have already stopped reading.

Reply ↓

Andrew on May 8, 2020 4:32 PM at 4:32 pm said:

Somebody:

One interesting twist here is that the U.K. has a Conservative Party government. Can you imagine how strong the reaction would be if all this were done by the Labour Party? One of the challenges for a conservative commenter such as Fund is that he is criticizing a conservative government here.

Reply ↓

Ben on May 8, 2020 4:56 PM at 4:56 pm said:

> One of the challenges for a conservative commenter such as Fund is that he is criticizing a conservative government here.

What's hard about that? I don't know John Fund's specific conservative stances, but there's plenty of in-party critiques in anything. It's standard American conservatism to critique government, and that goes on even when Trump is the president (https://www.businessinsider.com/trump-protestors-with-guns-in-michigan-capitol-are-very-good-people-2020-5).

And this doesn't come off as a critique of the government, it comes off as an accusation that scientists are somehow trying to mislead politicians by playing into their fears:

"I think politicians panicked in Britain – and some scientists enabled them – as they had in the past."

https://statmodeling.stat.columbia.edu/2020/05/08/so-the-real-scandal-is-why-did-anyone-ever-listen-to-this-guy/#comment-1332075

Reply ↓

Andrew on May 8, 2020 4:58 PM at 4:58 pm said:

Ben:

I'm not saying it's impossible for Fund to criticize a conservative government—he did it, after all!—just that it adds a degree of difficulty. If it were a liberal government, he could just go all out.

Reply ↓

Zhou Fang on May 8, 2020 5:04 PM at 5:04 pm said:

There's understood to be two factions within the UK tory party. There's the faction led by the Health minister, and then there's Dominic Raab and Iain Duncan Smith's faction who want the lockdown ended now.

Reply ↓

Jonathan (another one) on May 8, 2020 5:06 PM at 5:06 pm said:

We have the same problem here, Andrew, in which reflexive Trump defenders want to say Trump is perfect while saying Fauci is evil, when Fauci gives advice to Trump which he could just ignore! They seem blissfully comfortable with the dissonant idea that the infallible Trump is being suckered by Fauci, which would make him fallible, except he can't be. Bertrand Russell is smiling somewhere about the barbers who can't shave themselves in lockdown.

Reply ↓

confused on May 8, 2020 11:14 PM at 11:14 pm said:

Yeah, this aspect of the US discourse is a bit odd. (For that matter, Trump could simply fire Fauci and replace him with someone else; the Trump administration hasn't exactly been shy about turnover among its staff even at the highest levels!)

I think, though, the concept is that Trump has become so surrounded by "the establishment", "the swamp", or whatever that he's simply no longer getting any useful information on which to base decisions. (A lot of people believe that the officially-reported death numbers are simply fabricated, for example; that New York City simply recorded every death as COVID-caused.)

But, of course, this isn't how it really works; Trump is absolutely getting input from people who understate the crisis rather than overstate it.

Reply ↓

Bob on May 8, 2020 4:53 PM at 4:53 pm said:

3. I don't think the logic is totally unsound. I'll happily bet professional Microsoft coders introduced fewer bugs than academic coders who had their entire code base for the model in a single file. Don't forget, Ferguson himself said it was undocumented.

We all know his code is bad, that why he didn't release it.

4. That's an inaccurate representation. The issue was raised and Imperial said it was because of multi-threading, and would disappear on a single CPU. It didn't. Imperial's way of dealing with this was to average over many runs. This is fundamental, they introduced variation they couldn't explain. The person raising the issue saw a variation of 60k deaths.

"... also think that refusing to give out code on something as high stakes as this is pretty sketch, and that code quality standards in academia are pretty bad."

True.

Bear in mind though, this is the Gelman blog, populated by 'methodological terrorists'. Replicability is vital here, at least if you've got some rinky dink study about marshmallows or whatever. Instead, quite legitimate criticisms are waved away because a blog is 'agenda-driven'. Gee, the greatest world crisis since WW2 and a blogger has an agenda, unlike who exactly?

This is not a minor methodological point. The model risk here is enormous. It wouldn't be acceptable for a bank to value a basket of mortgages with such a model, the regulators would throw the book at you (today). It's fundamentally wrong that we have found ourselves making these decisions with these modelling standards.

Reply ↓

Zhou Fang on May 8, 2020 4:56 PM at 4:56 pm said:

You fundamentally misunderstand what the issue was.

Reply ↓

Bob on May 8, 2020 5:16 PM at 5:16 pm said:

No I don't. The blog post above references these issues:

https://github.com/mrc-ide/covid-sim/issues/116

https://github.com/mrc-ide/covid-sim/pull/121

"But Edinburgh came back and reported that – even in single-threaded mode – they still see the problem. So Imperial's understanding of the issue is wrong. Finally, Imperial admit there's a bug by referencing a code change they've made that fixes it. The explanation given is "It looks like historically the second pair of seeds had been used at this point, to make the runs identical regardless of how the network was made, but that this had been changed when seed-resetting was implemented". In other words, in the process of changing the model they made it non-replicable and never noticed."

"Why didn't they notice? Because their code is so deeply riddled with similar bugs and they struggled so much to fix them that they got into the habit of simply averaging the results of multiple runs to cover it up… and eventually this behaviour became normalised within the team."

Running simulations with the same random seed to recover an identical result, is so fundamental with code like this that it is astonishing they had this bug and didn't notice, then misunderstood what was causing it.

"Imperial are trying to have their cake and eat it. Reports of random results are dismissed with responses like "that's not a problem, just run it a lot of times and take the average", but at the same time, they're fixing such bugs when they find them. They know their code can't withstand scrutiny, so they hid it until professionals had a chance to fix it, but the damage from over a decade of amateur hobby programming is so extensive that even Microsoft were unable to make it run right."

I think this is an accurate representation of the quality of the codebase. It's not properly controlled or tested. Like 'somebody' above points out, this is the refactored code. The original is in a single file and has no doucumemnation. Maybe these issues are in the original, or maybe not.

But either they've introduced these bugs during translation, which as the blogger says could be because the original is so disorganised that Microsoft couldn't translate it cleanly. Or, the bugs were also in the original. It could be the original is bug free and clear as day, but then why not release it? Why get Microsoft to tidy it up in the first place?

Reply ↓

Zhou Fang on May 8, 2020 6:00 PM at 6:00 pm said:

"Why didn't they notice? Because their code is so deeply riddled with similar bugs and they struggled so much to fix them that they got into the habit of simply averaging the results of multiple runs to cover it up… and eventually this behaviour became normalised within the team."

This is nonsensical.

Reply ↓

somebody on May 8, 2020 6:05 PM at 6:05 pm said:

Let's look at the actual words written on the issue instead of just linking to it.

> But Edinburgh came back and reported that – even in single-threaded mode – they still see the problem. So Imperial's understanding of the issue is wrong

They never claimed that multithreading was the cause of the problem. They mentioned the multithreading, then asked the person to provide their command line parameters so they could debug. This is because they're not about to spend time troubleshooting an issue which may be expected behavior in the first place. Telling people possible things they may be doing wrong for a smoother trouble shooting experience is a completely standard practice on issue trackers for large projects, which should be abundantly clear to anyone who has ever done this kind of software engineering. They close the response with

> It would be helpful if you could give us the command lines you were using to demonstrate these differences, so that we can investigate further.

so they clearly weren't claiming to already know what was the cause, and that they intended to fix it. What you continue to repeat, over and over, is just a lie. They did not claim that multithreading was the cause of the problem, they did not dismiss the problem as not needing to be fixed or acceptable within a model averaging framework.

Reply ↓

Zhou Fang on May 8, 2020 6:14 PM at 6:14 pm said:

Let me actually explain the bug here:

The key part is this:

> "Simply removing the "/L:" argument causes the code to run correctly, albeit without the speed improvement of re-using the network.
> Note that this is not a random fluctuation but a constant effect: multiple sets of results all re-using the same network file (with no variation of parameters) will be identical, all with the same discrepancy from the original."

The way the model works is by simulating the spread of infection through a network of connectivities (basically who meets who). This involves random chance in multiple stages – in particular in the random generation of the graph, and the random progression of the virus. The generation of the network takes a substantial amount of time, so at some point someone decided to allow the model to be *restarted part way through the run*. The network can be generated, saved, and subsequent models re-run using the same network.

The issue the Edinburgh team uncovered was that if you run

Set seed
Generate Network
Save Network
Run spread

... you get a different result from

Set seed
Load saved network
Run spread.

If you rerun the program from scratch each time you get the same result. If you rerun the spread model from the same loaded network you also get the same result. So the model *actually does replicate*. What went wrong is that the coders missed that the generation of the network itself involved random numbers, and so change the state of the RNG seed, and that the correct thing was to keep the seed from before the generation of the network, making both ways of running the model consistent with each other. This was the fix.

In *both* processes the correct way to get the final output of the model is to re-run with multiple RNG seeds. These models use randomness at every point so the difference in the death toll figures isn't "unexplained", it's the exact uncertainty we are trying to capture.

Reply ↓

Zhou Fang on May 8, 2020 6:28 PM at 6:28 pm said:

Incidentally, the reason why they didn't see this issue is that if you look at the code they use to run the model (https://github.com/mrc-ide/covid-sim/blob/master/tests/regressiontest_UK_100th.py), they *do* not use the initial spread run. They generate the network first, discard the spread result, and load the saved network for all remaining runs. This means their results will be consistent each time.

Reply ↓

dhogaza on May 8, 2020 6:46 pm at 6:46 pm said:

Zhou

Pity you have to spend so much time debunking such obvious bullshit. Thanks for doing so.

Reply ↓

somebody on May 8, 2020 6:30 PM at 6:30 pm said:

> But either they've introduced these bugs during translation, which as the blogger says could be because the original is so disorganised that Microsoft couldn't translate it cleanly. Or, the bugs were also in the original. It could be the original is bug free and clear as day, but then why not release it? Why get Microsoft to tidy it up in the first place?

Again, you're demonstrating that you haven't paid very much attention. As Zhou Fang points out, the bug on issue 116 is in a new feature that lets you export the network graph to persistent storage and read it for another simulation run. The reason is that building these networks, typically with stuff like the Watts-Strogatz algorithm, is a pretty expensive part of the simulation. Not a huge problem if you're making one set of predictions to publish, but prohibitive if you want to play with parameters that don't affect interpersonal connectivity but may effect the final outputs. So the bugs weren't in the original AND also weren't "introduced in translation." They're new bugs for entirely new features. And introducing new features like this is why Microsoft and other engineers are working on it, to extend the capabilities and enable more people to iterate quickly on model development.

Again, I'm sure the code is bad as a guess, but I'm not about to pretend it's anything more than a guess. Pretending issues with new features must have also been issues in the original code doesn't make the guess more valid.

Reply ↓

somebody on May 8, 2020 6:44 PM at 6:44 pm said:

"But Edinburgh came back and reported that – even in single-threaded mode – they still see the problem. So Imperial's understanding of the issue is wrong. Finally, Imperial admit there's a bug by referencing a code change they've made that fixes it… Imperial are trying to have their cake and eat it. Reports of random results are dismissed with responses like "that's not a problem, just run it a lot of times and take the average", but at the same time, they're fixing such bugs when they find them"

This is just a mischaracterization of what's going on. Please read the actual issue yourself.

https://github.com/mrc-ide/covid-sim/issues/116#issuecomment-617304550

They bring up the multithreading as a possible cause only so that the person asking can rule it out. Telling people how to eliminate user error just in case is a standard practice from software engineers when managing any kind of large repo, especially a public one. They do not conclude by telling the user to just average, they conclude with "It would be helpful if you could give us the command lines you were using to demonstrate these differences, so that we can investigate further."

There, they are going to investigate further, they categorically did NOT dismiss the bug as acceptable when it was found. You are simply spreading false information, please stop.

Reply ↓

Anonymous on May 8, 2020 5:11 PM at 5:11 pm said:

> 4. That's an inaccurate representation. The issue was raised and Imperial said it was because of multi-threading, and would disappear on a single CPU. It didn't. Imperial's way of dealing with this was to average over many runs. This is fundamental, they introduced variation they couldn't explain. The person raising the issue saw a variation of 60k deaths.

This is just not true. Read the damn issue tracker.

1. They did not say it was because of multi-threading, they said that they were aware of indeterminacy which could result from multithreading so the user could eliminate that as the cause of the issue. The comment closes with "It would be helpful if you could give us the command lines you were using to demonstrate these differences, so that we can investigate further." So they did not claim to know why it was happening.

2. They did not at any point advance averaging over many runs as a way to deal with it. The blog post quoted Fergusson as saying “The model is stochastic. Multiple runs with different seeds should be undertaken to see average behaviour,” which is indeed correct behavior and roughly standard practice. Standard statistical procedures such as the bootstrap or jackknife, for example, are equivalent to “multiple runs with different seeds” where the seed determines which datapoints get drawn or left out. (In practice, I believe this is usually accomplished by taking sequential draws from an RNG with one seed, though the details of how intentional hardware randomness is generated escapes me).

3. Averaging is also certainly not still Imperial’s way of dealing with it because they fixed the bug.

Again, the people who are making the criticism seem to have either not read the linked issue, don’t understand software, don’t understand monte carlo methods, or both.

Reply ↓

somebody on May 8, 2020 5:59 PM at 5:59 pm said:

@Bob

If you’re a non-technical person who just has to go by the post on “lockdown skeptic”, it’s forgiveable, but please do not trust that source. First of all, the writer is clearly out of their depth when it comes to statistical programming and modeling. Issues range from dismissing the word stochastic as “a scientific-sounding word for random” to a complete misapplication of the phenomenon of feedback loops in machine learning to criticize computing R0 and using that quantity to compute other things downstream of it. There is nothing to suggest computing intermediate parameters and then using those to compute final outputs is bad modeling or somehow constitutes a feedback loop; that itself describes the entire practice of hierarchical modeling. If you want to rail against hierarchical models, you’d hardly find a less sympathetic place than here to do so. This problem of misunderstanding the technical aspects is on top of essentially lying about what was written on the issue, presumably hoping nobody would go so far as to actually click on the links.

> On a personal level, I’d go further and suggest that all academic epidemiology be defunded. This sort of work is best done by the insurance sector. Insurers employ modellers and data scientists, but also employ managers whose job is to decide whether a model is accurate enough for real world usage and professional software engineers to ensure model software is properly tested, understandable and so on. Academic efforts don’t have these people, and the results speak for themselves.

While I think there’s some value to this criticism, suggesting that “all academic epidemiology be defunded” requires a more complete argument than “this study was bad.” It’s also a pretty dubious position considering that the insurance sector cites work by academic epidemiology! This person is pretty confident the data scientists at insurance companies can handle it themselves, but I’m willing to be they would be pretty strongly against the idea. Anyways, I bring this bit up because it reeks of the sort of bitter hackernews software engineer type who thinks they shouldn’t have to know the definition of conditional expectations to be taken seriously on machine learning.

Reply ↓

dhogaza on May 8, 2020 6:34 PM at 6:34 pm said:

somebody:

“First of all, the writer is clearly out of their depth when it comes to statistical programming and modeling”

Oh, this is clearly obvious. WTF:

“Imperial advised Edinburgh that the problem goes away if you run the model in single-threaded mode, like they do. This means they suggest using only a single CPU core rather than the many cores that any video game would successfully use. For a simulation of a country, using only a single CPU core is obviously a dire problem – as far from supercomputing as you can get. Nonetheless, that’s how Imperial use the code: they know it breaks when they try to run it faster.”

This has got to be the one of the dumbest things written by someone who claims to be a software engineer I’ve ever read.

Bob, if you don’t see why this is one of the dumbest things written by someone who claims to be a software engineer, then I suspect that somebody’s guess is correct – you’re very likely not a technical person and are simply parroting what you’ve read on that site. You don’t really need to be an expert in statistical modeling, Monte Carlo methods, or the like to see how stupid this statement is.

I'm not an expert on those things, but just about any competent techie should have at least a wide angle grasp of how these things work. Determinism certainly makes code easier to debug, to validate changes, etc. But relaxing determinism to take full advantage of parallel computation is perfectly fine.

So far the only criticism from lockdownsceptic that I would agree with is that a single 15,000 line C file is poor practice. Ferguson probably didn't feel comfortable using tools like make that make managing a program broken into multiple files easier to manage. However, that does NOT prove that the C program itself was poorly structured or poorly written. There's evidence that it was machine translated from FORTRAN, given the history of scientific computing I'd say this is very likely. However, this doesn't mean that the FORTRAN code it was translated from was bad code. Mechanically translation from FORTRAN to C will lead to awkward-looking C code, but with a quality translation tool it won't lead to broken C code and awkward-looking code isn't necessarily inherently bad code.

I call bullshit.

Reply ↓

> Zhou Fang on May 8, 2020 6:41 PM at 6:41 pm said:
>
> A relevant link.
>
> https://stackoverflow.com/questions/3830347/can-a-multi-threaded-program-ever-be-deterministic
>
> Reply ↓
>
> > dhogaza on May 8, 2020 7:02 PM at 7:02 pm said:
> >
> > Plenty of MP programs produce deterministic results but that requires a good definition of what you need to be deterministic, and a hell of a lot of locking, which potentially slows things down a great deal.
> >
> > An example would be the implementation of different levels of transaction isolation in database software, where the user is allowed to trade off performance against various guarantees of determinism.
> >
> > That's determinism at a high level, though. At a lower level, for example the relative location of individual data rows within data files when you run the same set of transactions on two different MP machines, won't be guaranteed, at least not in database software meant for production use. That level of determinism is sacrificed for performance, while still guaranteeing the the specified transaction isolation semantics are preserved.
> >
> > So the fact that absolute determinism can't be guaranteed isn't of much importance in practice. For monte carlo simulations it's perfectly reasonable to sacrifice a certain degree of determinism for speed that doesn't effect the running of the algorithm (if you sacrifice determinism to the extent that one thread might use a variable before another thread has initialized it, you're screwed, of course, but that's not what's being talked about).
> >
> > Reply ↓

dhogaza on May 8, 2020 6:14 PM at 6:14 pm said:

Anonymous:

"Again, the people who are making the criticism seem to have either not read the linked issue, don't understand software, don't understand monte carlo methods, or both."

They CLEARLY do not understand Monte Carlo methods. Or pretend not to. Certainly don't understand software as well as they claim they do. Or pretend not to.

Bob says:

"The issue was raised and Imperial said it was because of multi-threading, and would disappear on a single CPU. It didn't."

Another answer to improve the chances you'll understand.

The bug report that was filed didn't mention whether they were running it in MP or SP mode. The response is simply pointing out that if they're running in MP mode, runs with identical data are expected to differ. In SP mode, not. They didn't say it would disappear run in SP mode. The respondent asked for the command line, obviously to see if it was run in MP mode or SP mode, i.e. to determine if it was a bug or not. Upon learning it was run in SP mode, it was understood to be a bug, and it was fixed two days later.

Is this what all this bullshit hinges on?

Reply ↓

Carlos Ungil on May 8, 2020 6:30 PM at 6:30 pm said:

I feel sorry for those who have to deal with this kind of stupidity:
https://github.com/mrc-ide/covid-sim/issues/175

At least there are also constructive contributions. I have no idea of how good or bad the code may be, and no interest in defending or attacking it, but the insistence on getting the original code seems a red herring. This is the reply one such request in the issue trscker:

"The code here is essentially the same functionally as that used for Report 9, and can be used to reproduce the results. The refactoring Microsoft and Github helped us with restructured and improved the layout of the code, with some documentation, to make it somewhat easier to scrutinise, but was written with regression tests against a reference result set to ensure changes in code structure did not change behaviour. We do not think it would be particularly helpful to release a second codebase which is functionally the same, but in one undocumented C file. Our refactor aims to make it easier (with some effort we acknowledge) for the open-source community to comment on the current live code, which is in use today.

"Seeing as you have quoted John Carmack's twitter, I hope that you find his comments encouraging when he writes, "it turned out that it fared a lot better going through the gauntlet of code analysis tools I hit it with than a lot of more modern code. There is something to be said for straightforward C code. Bugs were found and fixed, but generally in paths that weren't enabled or hit."

Reply ↓

dhogaza on May 8, 2020 6:38 PM at 6:38 pm said:

Whoever is having to deal with this stuff has my sympathy. Good grief.

Reply ↓

dhogaza on May 8, 2020 7:14 PM at 7:14 pm said:

Oh, man, then there are the self-appointed experts like MrHash in this issue thread:

https://github.com/mrc-ide/covid-sim/issues/178

It looks like the developers were having fun answering him.

Reply ↓

Winston on May 8, 2020 9:24 PM at 9:24 pm said:

Personally, speaking as some on who's done a bit of programming for insurance companies, I nearly did the nose trick with a cup of tea when I read "Sue Denim"'s suggestion that epidemiology should be handed over to the insurance sector. We should definitely also be using Blockchain though.

It's also amusing to hear people whining about lack of documentation and talking as if the method is something Ferguson pulled out of his backside when basic SIR goes back to 1927 – this seems to be a stochastic variant of a compartmental model – and there is a perfectly good docs folder with further links in the github repo here: https://github.com/mrc-ide/covid-sim/tree/master/docs

I suggest that, for their next project, "Sue Denim" and Bob could give us their critique of Carmack's code for Doom, as annotated in Fabien Sanglard's excellent "Black Book": https://fabiensanglard.net/gebbdoom/ In these trying times, we are all sorely in need of a good laugh.

Reply ↓

Bill Spight on May 8, 2020 5:43 PM at 5:43 pm said:

General comment that applies to a lot of these discussions.

I find the belief curious that says that a lot of people welcome or desire doomsday predictions. Since when has Cassandra become popular?

Reply ↓

Dalton on May 8, 2020 6:10 PM at 6:10 pm said:

According to the Social Security Administration, "Cassandra" peaked in popularity in 1990. I suppose it's another thing we can blame on the millenials.

https://www.ssa.gov/cgi-bin/babyname.cgi

Reply ↓

Anonymous on May 8, 2020 8:25 PM at 8:25 pm said:

The modern internet is a weird place. It's a place where people are paid to work for organizations that engage in infowars to create fake accounts and post ideas all over the place to give the impression of real controversy, sometimes even the same people post counter-ideas, to give the appearance of active controversy because controversy itself legitimizes fake arguments. If people are willing to argue over a thing, there must be some reasonable aspect to both sides of the thing right?

For the most part Gelman's blog has avoided this, but to the extent that it points to politically heavy-hitting policy controversies, you'd expect it to get infected with this problem as well.

Reply ↓

Thanatos Savehn on May 9, 2020 12:50 AM at 12:50 am said:

Awesome! It's cafeteria food fights like these that make me wish I'd gone into academia instead of law. Seriously.

Reply ↓

dhogaza on May 9, 2020 10:28 AM at 10:28 am said:

Oddly it's exactly why I chose NOT to go into academia! :)

Reply ↓

Sam S on May 9, 2020 4:57 AM at 4:57 am said:

John Fund's attacks on Neil Ferguson are just repeats of now-cliched attacks that have been parroted on twitter and in UK newspaper below-the-line for several weeks now.

I have no problem with arguing about policy or science. But this isn't an argument, it's just repeated sloganism. Just as sitting outside Neil Ferguson's house with a long-lens camera was not a good way to have an argument.

In fact it's worse than this. UK politicians are blaming all their decisions on "scientific advice", trying to absolve themselves of any responsibility. At the same time they are cherry-picking which scientists they listen to, and killing off the ones that they don't. (Boris Johnson's former employer, the telegraph, broke the Ferguson-affair story.)

Reply ↓

Anoneuoid on May 9, 2020 7:08 PM at 7:08 pm said:

> *In fact it's worse than this. UK politicians are blaming all their decisions on "scientific advice", trying to absolve themselves of any responsibility*

It would be awesome if at the end of this we get separation of science and state. Politics only brings down anything it pairs with, the other member of the pair never improves politics. Religion, Science, whatever.

Reply ↓

Astrid on May 9, 2020 10:46 PM at 10:46 pm said:

Hello, I was just reading with interest all your thoughts above – all very interesting. One thing that comes across is the need to appreciate nuance and how much the context matters. I am however a little reluctant to absolve Prof Ferguson on grounds of technicalities.
1) It seems that here Prof Ferguson is too simplistically judged as the statistician author of just another academic work. In basic science any report within the 95%CIs is legitimate, even if the range is wide. But would that be acceptable in medicine, where different standards of precision and communication are expected? Ultimately, we are judging the work of DOCTOR Ferguson, a clinical epidemiologist, bound by specific professional principles. It is a doctor's job to make sure their estimates are meaningful and they communicate their opinion effectively to the needing public. These are not just desirable qualities, but fundamentals of the medical profession. Would you accept wide 95%CIs from your doctor before making a decision on some surgical procedure? (well, you shouldn't!)
2) Here's a question about chance errors. In all unbiased research we expect some uncertainty around the central estimate. We expect the true effect to fall, at an equal probability, to the left or to the right of the central estimate. I am very familiar with all major public health crises in the UK and the US in the last three decades. Prof. Ferguson and his professional advise have been central to the UK government policy in pretty much all of them. And in all of them, his central estimates were way above of what the world really witnessed, even in absence of any meaningful action (bird flu). Were we to statistically test against the null hypothesis of chance, there will be a hint of something unidirectionally wrong in Prof Ferguson's estimates. In all objectivity, how likely are we to recommend a surgeon with a similar performance history to a family member?

Reply ↓

Martha (Smith) on May 10, 2020 1:21 AM at 1:21 am said:

Astrid said,

"In all unbiased research we expect some uncertainty around the central estimate. We expect the true effect to fall, at an equal probability, to the left or to the right of the central estimate."

Who is the "we"you are talking about? The expectation you describe sounds like it is what someone who doesn't have a strong statistics background might expect. Unfortunately, technical uses of words that are also used in non-technical settings often have different definitions from the non-technical uses. In particular, in cases where the scientifically appropriate model for some phenomenon is not a symmetric one, the estimate of a range of plausible values is in fact expected *not* to be symmetric around the point estimate. (Also, the words "biased" and "unbiased" have technical meanings that are not the same as their everyday meanings.)

Reply ↓

Dale Lehman on May 10, 2020 8:10 AM at 8:10 am said:

I agree with Martha's comment and I have another issue with what you say – and this does not mean I am defending Ferguson's past and present work. To state: "Would you accept wide 95% CIs from your doctor before making a decision on some surgical procedure? (well you shouldn't)" is mind boggling to me. If the true uncertainty is wide, then I would not accept my surgeon misrepresenting that fact by stating an unrealistically narrow CI. In fact, much of the criticism of the COVID models over the past weeks has been that their CIs are too narrow. Uncertainty is a fact of life and what we want is an accurate statement about how little we know – I certainly don't want someone (particularly my physicians) to artificially pretend their is more certainty than we really have. In fact, there are many problems in medicine that result from pressure (whether external or internal) on physicians to sound more certain than they should.

Reply ↓

Martha (Smith) on May 10, 2020 4:42 PM at 4:42 pm said:

Dale said,

" If the true uncertainty is wide, then I would not accept my surgeon misrepresenting that fact by stating an unrealistically narrow CI. In fact, much of the criticism of the COVID models over the past weeks has been that their CIs are too narrow. Uncertainty is a fact of life and what we want is an accurate statement about how little we know – I certainly don't want someone (particularly my physicians) to artificially pretend their is more certainty than we really have. In fact, there are many problems in medicine that result from pressure (whether external or internal) on physicians to sound more certain than they should."

Agreed. In fact (in my personal experience), physicians' excessive certainty (AKA overconfidence) can lead to iatrogenic negative side effects of treatment — side effects that may be worse than the original problem that the patient came for help with.

Reply ↓

Berend de Boer on May 10, 2020 11:16 PM at 11:16 pm said:

I think these comments would not be completed if it didn't contain a link to the second code review of Sue Denim. There are several incorrect claims in these comments claiming the original code did work just fine single-threaded (i.e. not producing random results). That is simply not true, and we can prove it.

But I have to say this model is a waste of time as the original 1927 equation basically gets the same numbers.

The Imperial code/model seems to be more a form of obfuscation, where a computer model is used to lend weight to an argument.

Reply ↓

Ben on May 11, 2020 9:36 AM at 9:36 am said:

I dunno, I don't really think the Sue critiques have much in them.

In the first post (https://lockdownsceptics.org/code-review-of-fergusons-model/) Sue outlines how there was a bug in the code, someone reported it, the team asked for a reproducible example, the person provided it, and the team fixed the bug. All that happened in the open.

It also really seems like the Professional Software Development shtick in these posts is just a facade:

> I ran the simulation three times with the code as of commit 030c350, with the default parameters, fixed seeds and configuration. A correct program would have yielded three identical outputs. For May 7th the max difference of the three runs was 46,266 deaths or around 1.5x the actual UK total so far.

That's not really enough information to see what Sue is doing. Like sure, there's a commit number. But what exactly are default parameters? Sue's critique would not hold up to Sue's critique. As someone who does academic software engineering (and I'm not the most careful person), that's not really a great bug report.

It really feels like someone trying to find stuff to be mad about. Lots of ranting about how software engineering is better elsewhere, and not much discussion in the way of the epidemiology here. Like, those are fine points that could be argued, but that's about a tweet's worth of contents.

If Sue is that concerned with bugs in the code, it seems like the team is fixing bugs and that's happening on Github, so Sue should go there.

If Sue is wanting to sue Imperial/Neil/the UK government for how the code was used in March, that's presumably something to take to court.

> A few people criticised the suggestion for epidemiology to be taken over by the insurance industry. They had insults ("mad", "insane", "adding 1 and 1 to get 11,000" etc) but no arguments, so they lose that debate by default. Whilst it wouldn't work in the UK where health insurance hardly matters, in most of the world insurers play a key part in evaluating relative health risks.

If Sue wants to privatize epidemiology? I don't know how to do that lol. Maybe a blog is the way.

Reply ↓

Zhou Fang on May 11, 2020 12:29 PM at 12:29 pm said:

The "default parameters", by which they mean the example code, uses multiple threads, so it shouldn't be expected to be deterministic. I don't really know why these deniers are going on about this non-determinism stuff given the devs have added a test for checking that the output is byte for byte identical. https://github.com/mrc-ide/covid-sim/blob/master/tests/regressiontest_UK_100th.py

Reply ↓

dhogaza on May 11, 2020 12:57 PM at 12:57 pm said:

Zhou Fang:

"I don't really know why these deniers are going on about this non-determinism stuff"

Because they are unable to show that the model doesn't properly implement the mathematics of the methods portion of Ferguson's flu epidemic paper that led to its development. Just like those who attacked NASA GISS's Model E when it was open sourced were unable to show that the model didn't properly implement the underlying physics as described in a series of published papers and the like. "FORTRAN! Not enough comments! Poorly structured!" was the battle cry.

The developers have answered the non-determinism issue in multi-processing mode:

Network creation is non-deterministic in the sense that individuals aren't guaranteed to be assigned to the same household (and possibly other minor variations). This seems to be because they don't make any effort to serialize the calling to the random function used to do the allocation, to make it run faster. However the various allocations are guaranteed to match the parameter values that are given that control what percentage of individuals are in multi-person households and the like.

Then we have the inconvenient fact for "Sue Denim" and the rest that the resulting network can be saved, and if loaded and run multiple times with the same seed values to drive the monte carlo simulation with the same number of threads on the same machine you'll get the same model results.

Reply ↓

dhogaza on May 11, 2020 1:03 PM at 1:03 pm said:

Also ...

"I ran the simulation three times with the code as of commit 030c350, with the default parameters, fixed seeds and configuration. A correct program would have yielded three identical outputs. For May 7th the max difference of the three runs was 46,266 deaths or around 1.5x the actual UK total so far."

Wanna bet that the default parameters are for the "no intervention, no voluntary social distancing" case? The case that yielded up to 500K deaths? Making "Sue Denim"'s comparison with actual deaths in the lockdown scenario dishonest?

Reply ↓

dhogaza on May 11, 2020 1:24 PM at 1:24 pm said:

Berend de Boer:

"But I have to say this model is a waste of time as the original 1927 equation basically gets the same numbers.

The Imperial code/model seems to be more a form of obfuscation, where a computer model is used to lend weight to an argument."

Which makes it clear that you, like all of the detractors thus far, haven't read Ferguson et al's original paper describing the model's purpose and goals.

The model's purpose was to investigate what level of early intervention would be required to wipe out a new and potentially deadly flu strain crossing from animals to people. How early do you need to catch it? Since tamiflu existed, how many doses would be needed to knock it down before it grew to epidemic proportions and then became pandemic? What role could social distancing play? Would closing schools alone be suficient? Stuff like that. How might such an epidemic spread geographically over time?

They used demographic information for Thailand, because at the time they state that the quality of data for that country was better than for any other southeast asian country, and southeast asia was a likely place for such a species crossover to take place.

That's useful shit, dude.

It seems apparent that the model sat idle and of course no such crossover happened in Thailand for them to calibrate against.

Go forward to earlier this year. It appears that in the UK at least there wasn't another model available to do this kind of modeling. The IC model itself wasn't restricted to the flu and Thailand, of course, that was parameterization and data. So they worked up the necessary data as best they could given what was known at the time.

As you said yourself, the worst case infection and death numbers give the same basic results for those two gross measures as does the model.

What it doesn't give you is an easy way to probe questions like "how effective would closing schools be?". So it was used to explore various scenarios, and the model outputs came with fairly wide uncertainty intervals. 7,000 to 20,000 deaths for the lockdown scenario, considerably lower than the actual deaths seen under this scenario. This isn't necessarily due to any structural issue with the model, because they used an R0 of 2.4 or 2.6 (I've seen both quoted), while the true R0 for the UK, at least, appears to be about 3. The initial R0 value makes a huge difference in the projections regardless of how sound the model is.

So how about you go read the paper before suggesting its purpose is to obfuscate? I'm not an epidemiologist and I found the paper easy enough to understand, though I didn't dig into the detailed methods section. You don't need to in order to understand the broad purpose of the model and how it works in a general way.

Reply ↓

Michael on May 11, 2020 5:04 PM at 5:04 pm said:

This is the real scandal

https://lockdownsceptics.org/code-review-of-fergusons-model/

Reply ↓

Zhou Fang on May 11, 2020 5:59 PM at 5:59 pm said:

Yeah, the fact that science denial sites like that exist.

Reply ↓

CHRIS EDWARDS on May 17, 2020 3:15 PM at 3:15 pm said:

This is all very intelligent comment, thank you to the contributors. I am a UK citizen Im not a coder nor a scientist nor a doctor. I think however there is a lack of base information. The following link to an article from 2005 sets out the origin of the codes used to model this pandemic.

https://www.newscientist.com/article/dn7787-flu-pandemic-lethal-yet-preventable/

So first the model assumption was flu not covid. Second it was modelled in a specific country where different norms of behaviour prevail than in UK/US, the elderly are in the community not in care homes.Third the actual behaviour of the flu modelled in the 2005 study is now known , but I am not aware this actual spread behaviour has been incorporated. Fifth it seems he assumed an equal spread of the virus and equal impact across age groups.

The critical point here is that most flu does impact fatally all groups including children and young adults, whereas some flus are better resisted by older people than younger people, the Spanish flu showed this.

There are many studies that look at the effect of various types of virus related illness on large groups and they show that elderly people have impaired immune systems and are more prone to illness and death than other groups for many types of illness and that corona virus is more prevalent than flu

https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7121031/

I would have thought it obvious that if the modelling shows
you have planned for a flu pandemic and the new pandemic invalidates your stockpile of medicine and vaccine
if not caught early enough the results will be devastating on medical services
that the elderly are going to be equally vulnerable

Hospitals have assumed they will need to care for all population groups and they have over provided beds , and when the most vulnerable ie the elderly present first they have been discharged from hospital to make way for the rest ot the population to care homes and locked in with no sunshine or exercise and low immunity, you have lit the match locked the doorh and thrown the tprch in.

So with Ferguson it really was gigo. Nobody challenged his methodology or assumptions, the code will reflect these.In the UK 90% of covid deaths are in the 65% age

https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/deathsinvolvingcovid19englandandwales/deathsoccurringinmarch2020

and the same report shows the peak age deaths formales aged 90+ age at 1725 per thousand whereas the flu in 1999 went above 2000 for all age groups on several days in December January and did not return back to 5 year average until April 2000

So – how serious actually is covid in the scale of events? Serious but it is an older persons illness primarily compared to pandemic flu which is generally non discriminatory, we used a flu omdel on a coronavirus surely a basic mistake and not one that a virologist would make; ferguson is not a virologist.

Reply ↓

Zhou Fang on May 17, 2020 4:11 PM at 4:11 pm said:

He didn't assume the things you think he assumed.

Reply ↓

dhogaza on May 17, 2020 4:35 PM at 4:35 pm said:

Chris Edwards

Zhou Fang gave a short answer but here's a bit more.

"So first the model assumption was flu not covid."

The model is parameterized with the characteristics of the disease being spread. If you're implying that characteristics like the basic reproduction rate was assumed to be the same for covid-19 and the flu, you're wrong. And you could've found that out yourself by digging a bit rather than (as I'm assuming you did) read something someone on the internet said that fit with what you wanted to hear.

"Second it was modelled in a specific country where different norms of behaviour prevail than in UK/US, the elderly are in the community not in care homes."

The model's demographics are also parameterized and data for the UK was used for modeling the spread of covid-19 in the UK. You could've found that out yourself.

Your third bit I find uninteresting, and you didn't specify a fifth bit.

"Fifth it seems he assumed an equal spread of the virus and equal impact across age groups."

It is equally infectious across age groups in equivalent settings. In his model, the elderly aren't placed in schools of course, so it does model differences due to things like this. I'm not sure if they incorporated information on nursing home and other long-care facilities. They started parameterizing the model for UK demographics and covid-19 fairly early on and the full dangers of the infection getting into nursing homes might not have even been really understood at that point. Red herring, anyway, as the model is being modified to deal with new kinds of demographic information.

"we used a flu omdel on a coronavirus surely a basic mistake and not one that a virologist would make"

To repeat, the model is a generalized agent-based epidemiological model that originally was *parameterized* for making projections regarding various flu epidemic scenarios.

That does not mean it is a "flu model". It incorporates a standard SIR model, for instance, and these are used to model all sorts of diseases.

Reply ↓

Luca Beltrame on May 28, 2020 2:18 AM at 2:18 am said:

I came to this specific post by chance and I see a lot of discussion on the IC model and on the IHME model. An example of a "model gone wrong" is however a less known story (or at least I think it is: after reading my comment, did anyone hear of it?) is what happened to Italy.

When discussing the country's reopening on May 4th (that wasn't a reopening, but I digress) the government expert panel released a scathing scenario (out of 92 being elaborated) that 150,000 Italians would have required ICU treatment by the end of the month. Of course this raised quite a few eyebrows, but it was defended vigorously in the media (happened on TV, so I don't have a link at hand), despite some concerns by other scientists.

The report, unlike IC's report 9, was not public. It was then "leaked", without any author information. Some people took a look at it and it was a bit scathing:

– IFR assumed at 0.657% with no explanation
– Overestimation of results (basing on the hardest hit areas from a month before the report)
– Outdated epidemiological models
– No correction for gender imbalance

I have links which I can provide for the critique, but you'd have to speak Italian to understand them.

IMO that was a spectacular blunder, because it eroded trust (trust that is also difficult to keep due to conflicting public statements from experts, e.g. there's someone supporting indefinite lockdowns until elimination that said that masks and social distancing were useless, while others have different stances).

In this case, better to have the models and the equations to look at. And for governments to release them, of course.

Reply ↓

Nigel Goodwin on May 29, 2020 2:04 PM at 2:04 pm said:

Can I get back to some statistics, although I am very grateful for the abundance of information here? My background is in probabilistic forecasting of complex computer models. Particularly for oil and gas reservoirs, where there is a reservoir simulator, and a simulation deck containing parameters for the particular model. The simulation may take from minutes to days to perform a single run. I use surrogate models and HMCMC (NUTS). The Covid code seems to have a similar structure, with similarly lengthy simulation runs. I did not see any evidence that the code has any kind of MCMC, nor defines any likelihood function, but I may not have looked hard enough, and it may be that others have written wrappers for thos kinds of calculations.

In the reservoir world, we have the problem of fitting to historical measurements, and then the problem of probabilistic forecasting. I have applied Bayesian techniques for both, where the number of parameters may vary up to 1000.

My question is – has this covid model been used for any kind of probabilistic forecasting? I see some kind of single parameter sensitivity analysis, but the oil industry got beyond that about 30 years ago.

If not, what collaborative possibilities are open.

How would you rate the current covid modelling in the context of bayesian forecasting?

I have seen attempts to define a pandemic ODE and use Stan, but from what I see this is not a feasible approach with this covid code.

Reply ↓

Andrew on May 29, 2020 2:12 PM at 2:12 pm said:

Nigel:

I've not looked into those old Imperial College models. When I've looked at Imperial College models, it's been the new stuff that is statistical and uses Stan; see for example here.

Reply ↓

John on May 31, 2020 2:59 PM at 2:59 pm said:

This was a fascinating read. I came here because I read an article on NPR about the difference in case fatality rates, which calculates the odds that someone who develops symptoms will die and infection fatality rate — the odds that an infected person will die. The headline reads. “Antibody Tests Point To Lower Death Rate For The Coronavirus Than First Thought” Was this a part of the original model?

Reply ↓

nigel goodwin on June 9, 2020 7:39 AM at 7:39 am said:

Been tearing my hair out at

https://www.linkedin.com/feed/update/urn:li:activity:6674508533352022016/

where they use the example of Rand.cpp at https://github.com/mrc-ide/covid-sim as an example of bad Ferguson code.

This wasn’t even written by ICL, but is a copy of code at RANDLIB which is used in the Python module Math::Random

All these arm waving dribbling attacks on the code, which do not make any reference to actual source files, have lost all credibility with me.

Reply ↓

Michael Carlotta on June 18, 2020 8:21 AM at 8:21 am said:

Often the elephant in the room is ignored.

Forget about flawed computer models. Computer models, good or bad, are always used by someone with a political agenda.

Take global warming models. The modelers and those blaming man ignore the VOSTOK Ice Core elephant in the room which shows 500,000 years of earth temperature and CO2 levels. The data shows we are in a predicted warming period, near the peak, headed for the next cooling cycle (relatively soon). The data also shows that CO2 increases follow global warming, is not the cause of. There are other elephants.

COVID-19 has its share of elephants. Mostly political and financial. It is not hard to see the COVID elephant hiding behind the floor lamp in the living room.

Reply ↓

Zhou Fang on June 18, 2020 8:50 AM at 8:50 am said:

https://skepticalscience.com/co2-lags-temperature.htm

Reply ↓

CROM on August 7, 2020 8:57 AM at 8:57 am said:

Zhou Fang, to invoke skepticalscience.com as a credible source is like bringing an active KGB agent to give you a clear view of the Soviet union. The site is owned and run by John Cook of the fraudulent “97% consensus” study, a study that was aiming to raise awareness in the public (Cook admitted) BEFORE EVEN looking at the papers which he supposedly used. He had an agenda and fixed conclusions below looking at his data. Moreover, true scientists that reviewed his paper found just 0.3% instead of the 97% consensus he claimed (Legates, Soon et al). Veer away from that propagandistic site.

Reply ↓

Zhou Fang on August 7, 2020 9:17 AM at 9:17 am said:

Lol “true scientists”.

Reply ↓

zdb on June 27, 2020 11:05 PM at 11:05 pm said:

Um. Sweden much? Considered a better studied disease and its protocols, like perhaps TB? What exactly is Ferguson’s field of study? Epidemiology? Odd that. https://en.wikipedia.org/wiki/Neil_Ferguson_(epidemiologist)

Wiki lists him as an epidemiologist. His academic career is all physics, theoretical physics, math and philosophy. A close friend died of AIDS so he flipped his computer models from “doctoral research investigated interpolations from crystalline to dynamically triangulated random surfaces” to modeling infectious diseases?

Someone here cuts Ferguson a break because he’s quoted as predicting 50 to 50k deaths over BSE when the media only published the upward bound. That seems odd as well. What is the point of such a prediction? It’s also accurate to predict that between 10 and 7.8 billion people will die in the next 50 years.

It misses the fact that Ferguson’s predictions and the subsequent destruction of small farmers in GB more likely caused more farmer suicides than deaths from BSE.

Reply ↓

CROM on August 7, 2020 8:43 AM at 8:43 am said:

I was going to comment about the prediction for BSE as being completely obsolete (as would – almost – a random estimate be). I agree completely.

Reply ↓

CROM on August 7, 2020 8:40 AM at 8:40 am said:

I can understand why scientists are trying to find excuses for works like his. It’s because of errors like his that science progresses, BUT, we should acknowledge the susceptibility of mathematical modeling to garbage-in garbage-out (GIGO). A all levels.
It is a human proclivity to interpret reality at a shallow level, with subpar quality evidence in hand and with arbitrary assumptions at play.

People should be more aware about the relative findings in the split-brain studies by Roger Sperry and Michael Gazzaniga. The latter named the left brain “the interpreter”. Iain McGilchrist’s work is extremely relative as well.

This is getting ridiculously dangerous for the progress and well-being of society.

Reply ↓

Priffe on September 27, 2020 12:54 PM at 12:54 pm said:

He predicted 80000 for Sweden without lockdown.
There was no lockdown. The Swedish agency checked the data and found it lacking, and did their own model.
They came up with 6-8000, worst case 20000.
We are at 5800 and will probably have 6000 by Xmas.
There is little spread here, noone’s wearing a mask and the ICUs are empty.

And Ferguson is an ass.

Reply ↓

Paul Wright on November 23, 2020 6:00 AM at 6:00 am said:

> He predicted 80000 for Sweden without lockdown.

No he didn’t. https://twitter.com/imperialcollege/status/1257991339364560898?lang=en clarifies.

Reply ↓

Selassie on December 16, 2020 5:29 PM at 5:29 pm said:

He claimed Sweden would pay a huge price for no lockdown, with 40,000 COVID deaths by May 1, and 100,000 by June. Yet in the first week of May Sweden only had 2,854 deaths.

https://www.nationalreview.com/corner/professor-lockdown-modeler-resigns-in-disgrace/

Reply ↓

FabianB on January 19, 2021 6:50 AM at 6:50 am said:

Selassie, you come here in December, skip the entire comment discussion (wich would have taught you more than everything your read on the topic in the entirety of 2020 apparently), repeat a claim looooooong understood as wrong and you cite as source.... JOHN FUND?

Reply ↓

Andrew on **January 19, 2021 8:25 AM at 8:25 am** said:

Paul:

I followed your link, and it reminds me of why I hate twitter. Here is the linked thread (from 6 May 2020) in its entirely:

> *THREAD: You may have seen false claims that Imperial COVID-19 "modelling envisaged Sweden paying a heavy price for its rejection of lockdown, with 40,000 Covid deaths by 1 May and almost 100,000 by June". Our researchers made no such prediction*
>
> *Professor Ferguson and the Imperial COVID-19 response team never estimated 40,000 or 100,000 Swedish deaths. Imperial's work is being conflated with that of an entirely separate group of researchers*
>
> *The Imperial team's models are published here http://ow.ly/sAW950zyleY and the source code for Report 13 on Europe is available on GitHub*

This is sooo twitter: a link, a statement, and nothing to back it up. I'm not saying these Imperial people are wrong in that post, just that they are offering no substantiation. If the Imperial team never made those predictions, it would help to (a) point to who falsely claimed they made those predictions, (b) where the 40,000/100,000 prediction actually came from, (c) what was that entirely separate group of researchers who made those predictions, and (d) what exactly did the Imperial team predict for Sweden in April. It's fine to link to their paper, but linking to their paper isn't enough. If you want to make an argument, you gotta fill in these gaps; assertion isn't enough.

Selassie:

The article you link to by Fund had lots of problems; see above discussion thread.

Reply ↓

FabianB on **January 19, 2021 6:45 AM at 6:45 am** said:

>He predicted 80000 for Sweden without lockdown.

you come into this threat in September 2020, when the Imperial College London explained in May 2020

"Professor Ferguson and the Imperial COVID-19 response team never estimated 40,000 or 100,000 Swedish deaths. Imperial's work is being conflated with that of an entirely separate group of researchers"

(Source: https://twitter.com/imperialcollege/status/1257991339364560898)

>There was no lockdown.

Really, that one didnt hold up either:

"The Swedish government, which has thus far left most schools, businesses and restaurants open during the pandemic, announced the closure of non-essential public workplaces, such as gyms, pools and libraries.

Prime Minister Lofven also said the wearing of face masks would now be recom-

mended on public transport during peak hours.

"Now we see that we need to do more because we see that the spread of the infection is too serious and we have a strained situation in the healthcare system still," Lofven said.

The government asked citizens to limit gatherings to eight people but there are no penalties for breaking the rules.

Lofven stopped short of imposing a "very serious lockdown" as it "wouldn't have an effect in the long run because people would not put up with that.""

(source: https://www.dw.com/en/coronavirus-sweden-implements-tougher-restrictions-as-cases-rise/a-55992637#:~:text=The%20Swedish%20government%2C%20which%20has,as%20gyms%2C%20pools%20and%20libraries.)

>And Ferguson is an ass.

Well at least he has adhered to standards of academic decency by not repeating long refuted claims.

Sweden, one of the least densely populated countrys in Europe had (despite, although to late to little measures) ~10.000 Deaths and a per capita rate in the league of places like the US or Bulgaria while its neighbours has a quarter to a 10th of that per capita death reate. Even the King of Sweden explained that sweden failed in its response.

Reply ↓

> Poster on February 16, 2021 4:16 AM at 4:16 am said:
>
> FabianB, your quotes about Sweden are absolutely incorrect. Sweden has not had a lockdown in your definition. The prime minister quote you made is incorrect. Don't attempt to use a "Sweden actually did have lockdowns" argument against anyone.
>
> The lockdowns have been full-on madness, and meanwhile Sweden has been essentially normal-normal. With that said Swedes socialized less this past year than normal due to a personal sense of caution.
>
> Corona has been a bad flu, lighter than the 1993 flu in Sweden, and about 90% lighter than the Spanish Flu of 1918.
>
> What we have seen in many countries such as the United Kingdom, is explosive biofascismo and neosovietism.
>
> Reply ↓
>
>> Carlos Ungil on February 16, 2021 8:29 AM at 8:29 am said:
>>
>> > With that said Swedes socialized less this past year than normal due to a personal sense of caution.
>>
>> It may be partly due to a personal sense of caution and partly due to the fact that no more than eight people may be present at a public gathering or public event. The police has the right to cancel or dissolve an event that has more than eight attendees.
>>
>> Reply ↓
>
>> Joshua on February 16, 2021 11:28 AM at 11:28 am said:
>>
>> Poster –
>>
>> Rather than taking the time to list all the ways that you crammed erroneous statements into a short comment, I'll just link this video – the Sweden relevant portion starts at about 12:15.
>>
>> https://youtu.be/v341VNPgL50
>>
>> And that video only begins to show the wrongness of much of the rightwing rhetoric about Sweden.
>>
>> Reply ↓

Ron Carlin on November 28, 2020 2:50 AM at 2:50 am said:

Occam's razor compels us to find the simplest explanation matching the outcomes. Is it not obvious that human minds,including our own, are not capable of rational thought during a panic? And when the panic has subsided, e.g. with foot and mouth, the circuits of the brain that could now do a post-mortem have been cauterised beyond repair, meaning we can't learn from our mistakes?

Reply ↓

Anoneuoid on January 19, 2021 10:46 AM at 10:46 am said:

It is now mid January and only about 50 covid patients have had vitamin c levels checked (all were deficient). The only RCT reported a promising 50% drop in mortality, that RCT data was known since March and still has not been followed up.

Likewise, it has been known that could patients immediately improve with HBOT since Feb. This makes perfect sense since during HBOT the oxygen saturation of the blood can go over 100%. Almost no effort has gone into that either.

But we did implement the most expensive and dangerous interventions of lockdowns, early intubation, etc. After doing that things are actually worse than before these interventions were started but lets keep listening to the same people that have never been right this entire time.

Reply ↓

Anoneuoid on January 19, 2021 10:59 AM at 10:59 am said:

Then there is the obvious protective effect of smoking, ignored since SARS-1 in 2003. What is going on with that?

Btw, there is pretty much guaranteed to be a SARS-2b or SARS-3 strain that comes out, probably triggering ADE in those with low levels of anti-S1 antibodies to the current SARS-2. This topic has gone largely ignored as well.

Reply ↓

Chris Bach on May 17, 2021 11:23 AM at 11:23 am said:

A lot of discussion in the UK right now as to easing lockdown. SAGE's model focused solely on the danger of meeting people outdoors and is still scaring the living daylights out of a population that has been in lockdown for almost six months (for which the slogan has been "Stay At Home, Protect The NHS, Save Lives". During all that time, the WHO has been running a campaign – "Meet People Outdoors". Now we have a good vaccine rollout and case rates plummeting, everyone should be jumping for joy – instead they've put a new model put suggesting new variants could kill another 100,000, so the government is poised to bring in yet another six month lockdown.

And this is why the British public are do confused, angry and weary. Taiwan has had no lockdown and just twelve deaths. SAGE kept us cooped up for six months*, during which seventy thousand people have died. SAGE thinks this is a success story because they had originally predicted half a million deaths in the UK. I say their model is do bad they should be facing at the very least a government enquiry, and at the extreme, face a court hearing to see if their crass incompetence and refusal to follow the examples of Taiwan etc (or even the WHO guidance NOT to stay indoors) is worthy of a Manslaughter charge.

*ust this lockdown – we've had various partial lockdowns before that.

Reply ↓

Joshua on May 18, 2021 10:01 AM at 10:01 am said:

Chris –

Before pressing manslaughter charges, you might consider enlarging your orientation towards confounding variables, and to make your referencing of "lockdowns" more nuanced.

For example, Taiwan was fully prepared to implement interventions such as shelter in place orders if it became necessary but largely avoided that necessity by focusing on interventions such as encouraging mask-wearing and intensively supporting testing and tracing and isolating identified cases. That difference complicates vague comparisons of "lockdowns" between Taiwan and other countries.

And in Taiwan, I think it's fair to say, the public has a more nuanced approach to evaluating the bi-directionality and interconnectedness of individual and societal freedoms and benefits, as opposed to many of the western countries such as the UK where the publics tend to be more self-centered in a somewhat obsessive focus on individual benefit as if it exists in some kind of zero sum relationship with societal benefit.

Reply ↓

Leave a Reply

Your email address will not be published.

Comment

Name

Email

Website

I'm not a robot

reCAPTCHA
Privacy - Terms

Post Comment

- Administrative
- Art
- Bayesian Statistics
- Causal Inference
- Decision Theory
- Economics
- Jobs
- Literature
- Miscellaneous Science
- Miscellaneous Statistics
- Multilevel Modeling
- Political Science
- Public Health
- Sociology
- Sports
- Stan
- Statistical computing
- Statistical graphics
- Teaching
- Zombies

Anoneuoid on The COVID wager: results are in

November 21, 2021 4:44 PM

Yes, it was confirmed at the top of this thread. Also, the first dose of Pfizer vaccine was reported to...

Anoneuoid on The COVID wager: results are in

November 21, 2021 4:38 PM

I did not look into that, the cdc site has moved (or removed?) the info on the individual models. It...

Anoneuoid on Fake drug studies

November 21, 2021 4:34 PM

Yep, it is very interesting to look into. Also that a medical intervention that killed only ~0.05% was considered a...

Joshua on Fake drug studies

November 21, 2021 3:52 PM

Andrew - That's funny if it's the same guy. I found his analysis very compelling. It will be interesting to...

Martin Modrák on Design choices and user experience in Stan and TensorFlow Probability

November 21, 2021 3:32 PM

Thanks for writing this! I've never used TFP, so I'm glad to learn how they do things... One thing I've...

Andrew on Fake drug studies

November 21, 2021 3:24 PM

Joshua: Wait---is that video by Greg Kellogg? I know him from high school!!

Andrew on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 3:13 PM

Typo fixed; thanks.

Andrew on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 3:13 PM

The: "Obsession" . . . that's putting it pretty strongly, no?

Andrew on Not so easy to estimate the effects of school mask requirements . . . how to think about this?

November 21, 2021 3:12 PM

Country: I agree. I see the appeal of a mask requirement and also the appeal of no mask requirement, as...

A Country Farmer on Not so easy to estimate the effects of school mask requirements . . . how to think about this?

November 21, 2021 2:43 PM

> At the policy level, I see the appeal of a mask requirement in part because it addresses many parents'...

Joshua on Fake drug studies

November 21, 2021 2:15 PM

Andrew - You probably aren't going to want to take the time to watch this guy's video - but I...

Kevin Dick on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 1:57 PM

Even if this were truly a randomized experiment, shouldn't you still control for pre-treatment differences? I guess I'm asking if...

the_real_tiddlydump on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 12:36 PM

Andrew, of your two obsessions (RDD and Hoover), this is your healthiest. Easily my favorite posts. Trained in Economics, here....

Daniel on Event frequencies and my dated MLB analogy

November 21, 2021 12:24 PM

I think this example is missing an essential feature. "Was this cancer caused by the chemical leak?" is at least...

Jackson Monroe on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 12:18 PM

It seems wrong to say that the forcing variable is "just another" pretreatment variable we can control for. It assigns...

paul alper on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 12:17 PM

Typo in title: "seen" should be "seem." On the other hand, am I missing something subtle? An in-joke? A nod...

Daniel Lakeland on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 11:46 AM

Agreed, you should definitely do more adjustments and maybe local estimation with something other than LOESS, but first... just plot...

andrew mascioli on Design choices and user experience in Stan and TensorFlow Probability

November 21, 2021 11:04 AM

One thing that I really appreciate about Stan is that its role as a standalone language helps bridge gaps between...

Thomas Passin on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 10:56 AM

I have a treasured dataset of gaussian noise added to a concave-upward quadratic. If you take only the central third...

Andrew on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 10:38 AM

Daniel, There's a literature on fitting these with local linear fits (of which loess is one approach, not necessarily the...

Anonymous on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 10:36 AM

Journal of Health Economics

Daniel Lakeland on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 10:28 AM

The right way to fit these is with LOESS and a smallish bandwidth. If something is there it will show...

Eric B Rasmusen on An economist and his lonely belief

November 21, 2021 9:56 AM

Two thoughts: (1) Whenever I've been in an economics seminar, and speaker answers a question saying, "That's the way everybody...

AndyW on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 9:52 AM

I know you are not into model comparisons/Bayes factors so much (if I understand correctly) -- but I've always thought...

Eric B Rasmusen on Just another day at the sausage factory . . . It's just funny how regression discontinuity analyses routinely produce these ridiculous graphs and the authors and journals don't even seem to notice.

November 21, 2021 9:47 AM

It's funny that this is listed with the tag Economics-- but quite appropriate, I have to admit. People like me,...

Andrew on The connection between junk statistics and the lack of appreciation for variation

November 21, 2021 8:23 AM

Wannabe: This question comes up from time to time. Here's the answer.

Niko on Design choices and user experience in Stan and TensorFlow Probability

November 21, 2021 7:47 AM

Hm, my understanding is very limited but I don't think this is accurate. I *think* the layers are used to...

Niko on Design choices and user experience in Stan and TensorFlow Probability

November 21, 2021 7:35 AM

Thanks both of you, I'll have a look!

Daniel Lakeland on Design choices and user experience in Stan and TensorFlow Probability

November 21, 2021 12:24 AM

Based on https://proceedings.neurips.cc/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf My impression is that a "neural ODE" is where instead of thinking of a neural network as...

Daniel Lakeland on Design choices and user experience in Stan and TensorFlow Probability

November 21, 2021 12:04 AM

It probably will be niche because it really appeals to people with a certain background, which is a combination of...

Daniel Lakeland on Design choices and user experience in Stan and TensorFlow Probability

November 20, 2021 11:58 PM

Julia is an interactive language just like Python or R, even though compiled. If you're talking about the "time to...

Wannabe Scientist on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 11:42 PM

I have a question. I work with public health data that governments actively use to decide where to allocate resources. ...

Anonymous on Not so easy to estimate the effects of school mask requirements . . . how to think about this?

November 20, 2021 11:32 PM

I don't know what arguments are there against masking. A large randomized trial in Bangladesh has shown that masks reduce...

Andrew on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 8:31 PM

Sean: I don't like my bubble to be burst, but I appreciate that you apologized for it!

Andrew on Not so easy to estimate the effects of school mask requirements . . . how to think about this?

November 20, 2021 8:04 PM

Ambarish: It's not just the "appearance" of safety. I'm talking about actual safety, as it seems pretty clear that masks...

Ambarish Chandra on Not so easy to estimate the effects of school mask requirements . . . how to think about this?

November 20, 2021 7:36 PM

Thanks for posting this Andrew. The link you posted is broken, here is the correct one: https://github.com/ambarishchandra/CDC-critique/blob/main/CDC_critique.pdf Comments welcome. As...

Sean on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 7:01 PM

"In the old days, we'd have blurry photos of the Loch Ness Monster, the shroud of Turin, and spoon benders...

John Williams on Not so easy to estimate the effects of school mask requirements . . . how to think about this?

November 20, 2021 6:58 PM

My impression is that there have been a lot of observational studies of masking, so just because one doesn't replicate...

Eric on Not so easy to estimate the effects of school mask requirements . . . how to think about this?

November 20, 2021 6:53 PM

The link seems to be malformed in the post but likely should be: https://github.com/ambarishchandra/CDC-critique/blob/main/CDC_critique.pdf

Joshua on Fake drug studies

November 20, 2021 6:49 PM

With Michigan experiencing high numbers of covid cases and record hospitalizations, I wrote this piece for my local newspaper reminding...

Bob76 on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 4:11 PM

Andrew wrote: If you're a hack, you can convince yourself that weak evidence is strong enough. Fooling yourself can be...

Dogen on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 3:48 PM

Could be, although that's not what I was thinking about. I was more circling around the idea that quantitative reasoning...

Mark Phariss on Design choices and user experience in Stan and TensorFlow Probability

November 20, 2021 3:43 PM

FWIW, there is an interpreter for Julia (https://julialang.org/blog/2019/03/debuggers/#julia interpreter) which should allow for a more Python-like experience in the iterative...

John Williams on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 2:58 PM

Dogen, I think your basic point is that we evolved as social animals. That has good aspects (e.g., ethics), but...

Howard Edwards on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 2:37 PM

Thanks for the reference to Tversky and Kahneman's original use of the phrase "Law of Small Numbers". In discussions with...

Joe on Design choices and user experience in Stan and TensorFlow Probability

November 20, 2021 2:13 PM

I keep thinking that I should look at Julia, but sometimes it reminds me of self-driving cars — it's always...

Dogen on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 2:08 PM

Jim, I agree with you wholeheartedly that "tribal identity" is an oversimplification. It would have been better if I'd written...

Niko on Design choices and user experience in Stan and TensorFlow Probability

November 20, 2021 1:40 PM

Charles: What's up with these "neural odes"? Can we do them in Stan? What are they even exactly, can you...

Joshua on The connection between junk statistics and the lack of appreciation for variation

November 20, 2021 1:08 PM

Andrew - I'm sure it's just a coincidence.

Niko on Design choices and user experience in Stan and TensorFlow Probability

November 20, 2021 1:06 PM

I really hate working iteratively with compiled languages. I actually think it's great to have one interpreted language (in my...