EDIT: 2020/06/29: I am no longer confident in this, due to a) demographic change in who gets infected b) possibility of new clinically significant mutations (that I did not have enough time to vet) and c) potential for fairly significant treatments (that I again did not have time to vet).

Note: I wrote this preliminary doc originally to gather my thoughts in one place and share it with other amateur forecasters. I am not a professional nor do I claim any such expertise.

Infection Fatality Ratio (IFR) has many meanings. For simplicity, I'm thinking of the operationalization here: https://pandemic.metaculus.com/questions/3755/what-will-be-the-ratio-of-fatalities-to-total-estimated-infections-for-covid-19-by-the-end-of-2020/

[1]My best guess is that the all-things considered global IFR by the end of this year will be ~1.1% (50% credible interval: .85%-1.6%), and it'd be about ~1% for people infected by the end of June.

Frankly I don't think <0.15% IFR is at all supported by the data, and even 0.2% or 0.3% is very dubious.

I understand that my median estimate is relatively high compared to published IFR estimates, quite high compared to Metaculus median (.8%), and very high compared to Foretold (.5%), so I'll try to justify why:

1. Meta-analysis prior
    a. An interesting meta-analysis of the current literature: https://www.medrxiv.org/content/10.1101/2020.05.03.20089854v1 The median is around 0.75%, with a 95% CI around 0.49% to 1.01%
2. Estimates of IFR keeps going up
    a. There's a bunch of meta-information that seems useful, including that studies that guess at IFR rates keep going up.
        i. For studies/preprints from April, 3/7 of the studies included in the meta-analysis above have a median above 1%
    b. Unlike the author of the meta-analysis above, I am not convinced (by the implicit assumption in his aggregation) that this is a statistical martingale.
        i. Actually I reread the meta-analysis and it's worse than that: earlier papers (with lower IFRs) are given equal weight to later papers.
3. Strength of data is not regionally independent of true death rates, a common assumption
    a. People use studies on populations with very good data to extrapolate to other regions, and seem to assume that those regions and populations have IFRs representative of their country or even WHO region(!). (Or more charitably, high variance but unbiased samples).

[1] Which, incidentally, is plausibly worth reading in detail. A lot of my ideas for how to think about IFR came from the comments there (even though my final result ended up being broadly higher than most of the other commentators).

b. If *I* were to write a meta-analysis, I'd probably title it "*Early coronavirus cases and deaths are concentrated into regions with high transmission rates, good record-keeping, overall healthier populations, and a well-developed medical infrastructure.*" (as a homage to this excellent takedown of supercentenarian literature)
c. Some examples:
    i. Wuhan has a life expectancy of 81.1, higher than the US (78.5) and 4.6 years (!) higher than the rest of China: https://en.wikipedia.org/wiki/List_of_cities_in_China_by_life_expectancy
    ii. Lombardy has much better health infrastructure than southern Italy.
    iii. NYC has a median age 2 years younger than the US average, and life expectancy 2 years higher than the rest of the US[2].
d. I don't think this is random or pure good/bad luck, I think we'd expect better data collection in regions with better health infrastructure overall, and for city residents to be younger, generally healthier, and to transmit infectious disease[3] earlier and more quickly than people living in smaller cities or away from cities.

4. Statistical controls seems fairly one-sided:
    a. Research and estimates that try to control for biases in extrapolation seem to usually only control for biases in one direction. So they'd control for people in Lombardy and Wuhan being older than Sub-Saharan Africans, but not for people in Wuhan and Lombardy being a lot healthier (life expectancy of 54 in Nigeria, 81 in Wuhan, 83.2 in Italy. Can't find it for Lombardy but as noted above, should be higher in Lombardy than the rest of Italy).
    b. For example, this Imperial College report about IFR in developing countries *explicitly says* that they didn't control for populations in other regions being less healthy and have worse health infrastructure than China[4] so their estimates might be too low for low-income countries, but I still see people uncritically cite their purely age-adjusted estimates of ~4x difference between the UK and Subsaharan Africa.
    c. Also when people do forecasts around the presumed medians of things (eg assume that symptom onset to death is 12-20 days), they systematically under-estimate the long tail of people slowly dying from covid.

---

[2] To get an intuition of what this means at a slightly lower level, NYC has an adult obesity rate of 22%, which is high, but considerably lower than the US as a whole of 39.8% (!) obesity.

[3] Both for the obvious reason of interacting with more people, and because cosmopolitanism means they see international diseases first

[4] "To estimate the demand for health services and overall mortality, we use age-specific estimates of the hospitalisation rate and infection fatality ratio (IFR) obtained from our previous analysis of data from China4 . Hence, we make the strong assumption that similar levels of medical care to that provided in China are available elsewhere. We also implicitly assume that mortality patterns do not vary given the different co-morbidities. These assumptions may mean that our results may overestimate mortality in some HICs and under-estimate it in some lower income countries. " (Page 15). I actually slightly disagree with them about HICs since as noted above, I don't see a principled reason to use the baseline health infrastructure or population health of China rather than Wuhan, which has a life expectancy that's higher than the US, similar to the UK, and very slightly higher than the EU.

5. Systematically over-estimating age-structure effects
    a. I think in extrapolative forecasts, people *still* vastly overestimate age-structure (and gender) effects. Aside from really young people (which is primarily a factor in the under-five mortality rate, which is high for all cause mortality in some parts of the world but ~0 for covid), mortality rate from coronavirus per age group seems approximately proportional to all-cause mortality rate overall: https://medium.com/wintoncentre/does-covid-raise-everyones-relative-risk-of-dying-by-a-similar-amount-more-evidence-e7d30abf6821
        i. Certainly, being old is really bad for covid fatality, but so is everything else!
    b. I've been persuaded by Metaculus to use ratios of crude death rate overall as a rough proxy rather than age structure.
6. A crude estimate from NYC data
    a. On that note, I have *relatively* narrow error bars around ~.95% COVID-19 mortality rate in New York City. (first-pass guess ~1%, goes down a little for epistemic humility)
    b. Crude death rate from all causes in NYC is ~.546% in 2017.
    c. Crude death rate worldwide is .77%. If we subtract out under-5 mortality (eg, divide the rate by 4 and multiply by % of world population under 5), I think a rough first order guess is that the CDR for people over 5 worldwide is .67%.
        i. Note that there is a lot of country-level and regional variance, for example CDR in UAE is 0.143% (not a typo) and CDR in Ukraine is 1.45%, a 10x difference.
    d. So my first-order prior is that if everybody in the world is uniformly infected, the IFR would be 22% higher than the rate in New York City, whereas other people assume that the IFR would be on average lower in other places than NYC.
7. Media asymmetry, where media likes to report on low IFR estimates and places
    a. I worry that poorly done reports of lower IFRs are much more well-represented in the media and thus unconsciously bias people when doing future studies and forecasts. One obvious example: back when South Korea's naive CFR was like .6%, everybody talked about South Korea fatalities. Now that South Korea's well past the peak of their outbreak and their naive CFR is >2% (as far as I can tell, they haven't gotten worse at testing or have gotten a new outbreak, it's just that the death rate of previously confirmed cases slowly creep up), I still see reports of how Korea successfully controls their infection rate but it seems like nobody wants to talk about Korea CFR anymore. Instead all the talk these days is about Iceland?
    b. Likewise, really crappy studies of low IFR, like the Stanford serology studies and the Oxford CEBM estimates, gets cited a lot in the media, but never once has a really crappy study of >5% IFR ever passed my filter bubble
    c. I don't think updating on bad studies/media cherrypicking is a direct mistake researchers make, but it seems like a reasonable source of indirect bias.
8. Epistemic Humility + Conclusion

    a. I think the balance of factors so far suggest ~1.25% IFR. Since I have a lot of fundamental uncertainty and general estimates I've seen from other people are lower, my all-things-considered best guess IFR is 1.1%, with fairly wide tails.

Cruxes I have:

1. Treatment. Right now I'm imagining there's no likely future treatment or cocktail of treatments that a) can cut IFR by more than 1/3 of existing levels and b) will be widely used by most of the people infected by the end of 2020. If this assumption doesn't hold, my numbers can be way off.
2. Degree of badness of hospital overloading (or places that essentially don't have access to large-scale hospital care, period). We now know that the benefits of invasive ventilation are relatively minimal. However I have a bunch of uncertainty about the relative risk ratios of other things like noninvasive ventilation etc, so good studies on those will update me towards believing in more/less differences between NYC and middle-income or low-income countries.
    a. Not sure which direction, my *current* guess is that if hospital overloading contributes a lot to IFR, this would mean I currently under-estimate IFR in older middle-income countries and over-estimate it in younger low-income countries
3. Clarity about which regions are most likely to get infected
    a. Countries with younger and healthier populations (especially Middle Eastern oil countries) have something like 1/10 the crude death rate of older and less healthy populations (especially East European ones)
4. I'm currently assuming that the baseline probability of mutations that might change fatality rates in the last 5 or next 8 months is low.
5. I'm assuming that expected regional variation in viral load isn't very high. I can see a plausible case for viral load either not being a factor at all or affecting mortality by 2-4x in ideal conditions, but I just don't currently see a compelling case for it being a large factor in global IFR estimates.