Modelling projections for COVID-19 epidemic in *[Country X]*

LSHTM CMMID COVID-19 Working Group

Version 3: 16 June 2020

Aide for report interpretation

Background

On 5 May 2020, the Centre for Mathematical Modelling of Infectious Diseases (CMMID) at the London School of Hygiene and Tropical Medicine (LSHTM) published country-specific, standardised reports presenting projections of the COVID-19 epidemic under different response scenarios. All countries in Africa, Latin America and the Caribbean, South Asia and the Middle East are currently covered by the reports. The reports are downloadable from https://cmmid.github.io/topics/covid19/LMIC-projection-reports.html, and have since been updated to the above-dated version. Countries are listed by their three-letter acronyms. PDF reports are complemented by figures and output datasets.

The present document aims to guide users through the reports to facilitate correct interpretation and application of the estimates presented. Note that this aide refers to the above-dated version of the country reports, and will be superseded by later versions once the reports are updated.

What the reports cover

Geographical scope and time period

Each country report presents estimates for the entire country, over a 12-month period. The start date of this period is not explicitly specified in the report, but rather is a given "time 0" that corresponds roughly to the time point in each country when local transmission (i.e. sustained chains of transmission among people whose exposure to the virus was not foreign travel) would have been firmly established. The mathematical model used to develop the estimates simulates this by 'seeding' the country with 50 cases all at once (i.e. the model equations start with 50 infections at time 0).

As such, the start of the 12-month period in the reports is not quite the same as when the first case was confirmed in the country; rather, it is perhaps a few (2-3?) weeks later. Users should roughly work out when this time point is (e.g. if the first case was recognised on 15 March, the model projections would cover a time horizon from the first week of April onwards). At the same time, in some countries it is possible that the first infection was introduced before the first recognised, test-confirmed case: this can be determined with some approximation if pre-existing patient samples are available, and if sophisticated genetic analysis is done. If it becomes clear from the evidence that an earlier-than-recognised introduction occurred, the clock should be set back accordingly when considering these model projections.

Structure of each report

The country reports are completely standardised from one country to the next. Each report is comprised of the following identical sections:

1. **Overview**

   This section introduces the report's methods and explains / defines some of the quantities presented.

2. **Unmitigated COVID-19 Epidemic Trajectory**

   This section presents text, tables and figures describing the evolution of various quantities (cases, deaths: see below) in a scenario featuring zero control interventions of any kind, i.e. something akin to how the epidemic would progress within the population if nothing were done. This scenario is admittedly unlikely for most if not all settings but is useful to at least establish a counterfactual for how severe the epidemic *could* be in the country, against which to compare different response options. Quantities are presented for both the peak of the epidemic (i.e. the day with the highest predicted caseload) and cumulatively (i.e. since the start), and for different time points within the period.

3. **Intervention Scenarios**

   This section presents estimates of the percent (i.e. relative) reduction in different quantities (cases, deaths) that would be achieved by combinations of different interventions (detailed below), and quantifies how the epidemic would evolve under each intervention scenario, in terms of its timespan, the timing of the expected peak, and other quantities (cases and deaths), both at the peak of the epidemic and cumulatively.

4. **Methods, data and assumptions**

   This section provides summarised information about the methods and key data inputs, including specific data assumptions made for the country in question. It also acknowledges contributors and sources of funding that supported this work.

Information presented by the reports

Cumulative forecasts

The country reports provide various forecasts under the unmitigated scenario, and for the different intervention scenarios. These forecasts include the **cumulative number** (that is to say, since the start of the analysis period) of various quantities, overall and by age group, at different time points in the analysis period (3, 6, 9 and 12 months). These quantities include:

- symptomatic cases, meaning any SARS-CoV-2 infections that lead to symptoms of COVID-19; for any country epidemic, these will mainly consist of mild and moderate cases, with severe cases as a minority of the total;
- hospital bed-days due to COVID-19 cases requiring hospitalisation (both critical and non-critical), denoted in the report as 'Hospital Person-Days'. Bed-days mean the number of days that need to be spent in hospital by all patients combined (for example, if there are 3 patients in total and they spend 5, 17 and 11 days in hospital respectively, bed-days or hospital person-days = 33). It is a familiar quantity to hospital administrators, and useful for health service planning. This quantity is not a prediction of how many people will in fact be hospitalised, but rather of the demand in terms of bed-days that the country would need to meet: the model does not make assumptions about what proportion of this demand will be met.

- hospital bed-days due to COVID-19 cases requiring non-critical hospital care (i.e. non-invasive respiratory support, such as oxygen), denoted as 'Non-critical Person-Days', and thus a subset of the above;
- hospital bed-days due to COVID-19 cases requiring critical care (i.e. ventilation, renal support), denoted as 'Critical Person-days, as above;
- deaths directly due to SARS-CoV-2 infection; note that indirect deaths due to livelihoods and health service disruptions are not estimated by this model.

Forecasts at the epidemic peak

In addition, the country reports present forecast values of certain quantities **at the epidemic peak**, i.e. on the day that the highest number of new cases is predicted to occur:

- symptomatic cases (both severe and non-severe) newly occurring on the epidemic peak day (Incidence of Symptomatic Cases);
- deaths newly occurring on the epidemic peak day ('Incidence of Deaths');
- total cases requiring hospitalisation (both critical and non-critical) denoted in the report variably as 'All Hospital Demand' or 'All Hospital Occupancy'. More specifically, this refers to the number of patients that would need to be in hospital on the day of the epidemic peak in order to meet demand (or, alternatively, the number of hospital beds needed at the peak). It's a measure of prevalence (cases at a given point) rather than incidence (new cases per unit time).
- total cases requiring critical care (i.e. ventilation, renal support), denoted as 'Critical Care Demand', and being a subset of the above;
- total cases requiring non-critical hospital care (i.e. non-invasive respiratory support, such as oxygen), denoted as 'General Hospital Demand', and being a subset of the above;
- the estimated day of the peak, as a function of the start of the analysis period ('Peak Day'). Thus, '110' means 110 days after time 0, as described above.

Uncertainty ranges and their interpretation

All of the estimates are subject to assumptions (see below), and in addition incorporate uncertainty about some of the key parameters (variables) that underlie the model. The model is said to be 'stochastic', a technical term that simply means that it explicitly factors in known uncertainty and variability concerning various characteristics of the virus and its interaction with people: in short, it allows chance to play a role. For example, the duration of the COVID-19 symptomatic period is not a fixed quantity for everyone in the population, but rather shows considerable variability from one person to the next: the model represents this variability through a statistical distribution, itself based on observations made by epidemiologists carrying out surveillance, and published in the scientific literature. Indeed, most of the parameters that go into the model are statistical distributions, rather than fixed values. The model is run a very large number of times, and in each run a value is chosen at random from the distributions.

What one gets at the end is accordingly not a single forecast, but rather a range of possible forecasts for each quantity: this range is itself shaped like a distribution. Rather than presenting the most likely value (say, the average of the distribution of model runs), the report's tables present the following uncertainty ranges:

- The inter-quartile range, or IQR: this consists of the 25^{th} and 75^{th} percentile of the output distribution, i.e. the values in the output distribution that, respectively, 25% and 75% of the model runs fall below; think of this as the most likely range.

- The 95% interval (in parenthesis after the IQR): this consists of the 2.5th and 97.5th percentiles, and is perhaps more familiar to readers, as it resembles a 95% confidence interval. Think of this as a more conservative range, representing just how different reality could be.

It's important to note that uncertainty ranges of this kind don't mean that every value within the range is equally likely: values close to the average are more likely. On the other hand, it's also worth bearing in mind that these ranges capture only the known sources of uncertainty and variability. They give some idea of how much imprecision is associated with the forecasts, but they don't necessarily mean the forecasts are 'right': if the model's assumptions or input values are inappropriate and don't represent reality in the country being modelled (as discussed below), the forecasts could be considerably biased, in addition to being imprecise. The following diagram may help to conceptualise the difference between imprecision and bias (think of each black dot as a single stochastic model run):

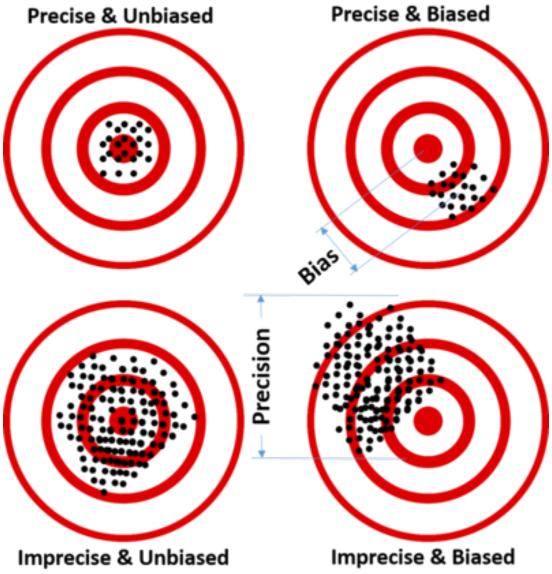

Source: University of Michigan School of Nursing (https://wiki.socr.umich.edu/index.php/SMHS_BiasPrecision)

Very imprecise forecasts are of course not very helpful, as it's difficult to take confident decisions when the ranges of imprecision are so wide. However, bias is potentially more harmful to decision-making, as it's not immediately visible.

Interventions considered in the reports

The country reports only consider non-pharmaceutical interventions (i.e. not requiring drugs or vaccines) based on different distancing measures that affect the amount of close, physical contact among people both within and outside households. The following interventions are featured in the reports, alone or in combination:

Intervention	Details	Scenarios considered	Notes
Distancing	Also known as general or 'social' distancing, this intervention applies to the entire population equally: everyone's contacts outside the household are reduced by a certain percentage. The reduction can be achieved in different ways, such as through behaviour change, curtailment of movements and gatherings, or improved hygiene. The model does not look at these specific measures, mainly because they are not really easy to	20% reduction in contacts. 60% reduction.	Note that reductions of 60-80% are consistent with lockdown conditions, i.e. very restrictive measures that would likely have serious economic and social costs. On the other hand, 20% might well be achievable

Intervention	Details	Scenarios considered	Notes
	tease apart mathematically. Think of this instead as the sum total of everything one could do to reduce physical contact among any two average people out there.		and sustainable without incurring such costs.
Lockdown	Distancing as above, but for a specified duration, and featuring a large reduction in contacts.	50% reduction in contacts over the first 30, 60 or 90 days since establishment of local transmission	
Intermittent lockdown	As above, but implemented repeatedly at regular intervals.	One week with 40% reduction in contacts followed by one week without any reduction, and so on	May be difficult to implement in practice.
Public health social measures (PHSM)	Term used by the World Health Organization. Here, this comprises of two interventions together: (i) general distancing as above, plus (ii) self-isolation of people with symptoms of COVID-19, for as long as they remain symptomatic: this is modelled as a reduction in contacts outside the household (the model assumes that it would be very difficult for sick people to also isolate from their household members).	20% reduction in contacts 25% additional reduction in contacts among symptomatic people	By its nature self-isolation cannot prevent transmission from individuals who are asymptomatic (i.e. never develop symptoms at all) or pre-symptomatic (i.e. already infectious but not yet ill). As such, it has limited (but not negligible) potential for impact.
School closure	Schools are closed, such that children and school staff have none of the contacts they would usually at school: the contact matrices used break down contacts by household, school, work and other settings: the model merely turns off all of the school contacts. The model does not compensate for school closures by assuming greater contact of children with their household members or neighbours.	0% school contacts	Consider the indirect harms of school closure, including worse educational, health and nutritional outcomes and child protection concerns.
In-home elder shielding	People at high risk of severe COVID-19 disease and death are supported to reduce their contacts both outside and within the household. In this modality of shielding, people are isolated individually within their household.		High-risk population groups are known to include the elderly and/or those with known diagnoses of certain co-morbidities (e.g. non-communicable diseases such as diabetes or cardiovascular disease; immuno-suppressive conditions). However, the model only considers age (60+ years) as the single criterion for being at high risk, i.e. eligible for shielding, and as such is likely to underestimate the true effect of this intervention.
'Green zone' (GZ) shielding	As above. However, in this modality of shielding high-risk people are grouped together, e.g. in designated shielding residences. The country reports assume that shielded people come into contact with other shielded people at the same rate as before shielding (in practice, this implies an assumption that people do shield together somehow, since individual shielding would reduce this contact to zero): this assumption is not explicitly stated in the reports. Furthermore, the reports assume varying levels of shielding uptake or coverage (C) and reduction (R) in contacts with unshielded people.	40C-80R, i.e. 40% of 60+ year olds are shielded and reduce their contacts by 80% 60C-60R, as above. 80C-80R, as above.	

The model assumes that the interventions (other than lockdowns) will be maintained for the entire period of 12 months; under most interventions considered by the reports, however, epidemics would not last this long (this can be seen in the reports themselves), and as such the actual period over which the interventions would need to be implemented could be considerably shorter.

Model assumptions and limitations

Brief description of the model

In epidemiology, mathematical models attempt to successfully simulate how epidemics progress. Most such models do so through a set of mathematical equations. Respiratory pathogens such as SARS-CoV-2 are relatively straightforward to model compared to infections with more complex mechanisms of transmission (e.g. malaria, which involves a mosquito, or HIV, for which many transmission mechanisms exist).

The model underlying the country reports is known as a Susceptible-Exposed-Infectious-Removed (SEIR) model: these terms refer to the 'classes' or 'compartments' that the model divides the population into. SEIR models are ubiquitous in studies of COVID-19 and other respiratory pathogens. The model underlying these country reports has successfully predicted the evolution of the epidemic in the UK and other high-income countries, and its structure is very similar to those of models implemented by other leading centres of excellence (e.g. Imperial College, Johns Hopkins University).

The model can be thought of as simulating two processes:

- Susceptible (i.e. not previously infected) people becoming infected: this process is driven by a few key parameters, including how *potentially* transmissible the virus is (we quantify this through the so-called basic reproduction number R_0, namely how many infections on average result from every case in a fully susceptible population); how much time on average elapses between different generations (i.e. chains) of transmission (the so-called 'serial interval'); and of course what interventions are actually being implemented: these include distancing measures that reduce the rate of physical contact among people, as detailed above; and vaccination, which makes people less or not at all susceptible. Both types of interventions reduce the virus' transmissibility.
- Infected people progressing to either symptomatic or asymptomatic infections, and some symptomatic cases becoming severe, critical or dying: these sequential clinical processes occur based on per-capita risks of different outcomes, including the risk of becoming severe if one is symptomatic, and the 'case-fatality ratio' (CFR, i.e. the proportion of cases who die) once cases become severe.

The model is also 'age-structured', meaning that it divides the population into different age groups, and, where relevant, applies parameter values that are specific to each age (for example, the risk of becoming a severe case is a lot higher as age increases). One additional, and quite important age-specific data input is the assumed age-specific contact matrix, meaning the extent to which different age groups, on average, have physical contact with each other: this is particularly relevant for COVID-19, for example because greater or lower contact between elderly people and children might considerably alter the former's risk of infection, or the impact of school closures.

Model strengths

As mentioned, SEIR models have proven validity and longstanding applications for predicting epidemic evolution and formulating public health policies. In addition to this general validity, the model underlying the country reports does attempt to some extent to reflect the pre-COVID-19 reality of each country. This is done by:

- Inputting the population size and age distribution of the country itself;
- Not assuming any particular R_0 value, but rather exploring a range of possible values, making the uncertainty ranges more likely to contain reality;
- Allowing for some increase in severity risk and CFR by age in Africa and the Middle East, compared to China and Europe: this is done somewhat crudely by (i) shifting the age distributions of risk downwards by 10 years (i.e. the per-capita risk of severe disease of a 70 year-old in China would be the same as that of a 60-year old in Iraq): this may capture some of the greater vulnerability and older biological age that characterises low-income settings; and (ii) multiplying CFR by a factor of 1.5, recognising that cases might not have much access to appropriate treatment.

Model limitations

Numerous caveats need to be borne in mind when interpreting the country reports. The two main limitations are listed here, and further points are tackled in the FAQ section below:

- The reports present hypothetical scenarios of what might happen. They are very naïve to what the actual potentially transmissibility (R_0) of the virus is in country, and do not integrate any surveillance findings from the country itself; equally, they do not take into account the effect of interventions being implemented in the country (e.g. lockdowns, test-and-trace, etc.). As such, they are best considered didactic tools to understand the plausible scale of the unfolding problem, and how well different interventions might work. The relative effect of different response options is relatively unaffected by model assumptions on transmissibility.

  Likely direction of bias: difficult to gauge. However, early interventions such as lockdown and aggressive testing and contact tracing of cases would probably have reduced transmissibility considerably, i.e. the country report might predict a much earlier epidemic peak than could in fact occur.

- Some of the key parameters of transmission and disease severity may not be reflective of the reality within the country. To date, the best estimates of R_0 , the serial interval, the duration of different stages of the infection and disease, as well as the risk of (severe) disease as a function of age, come from China, Europe and the Diamond Princess cruise ship outbreak. As discussed above, the country in question might in fact experience a higher or lower R_0 than that seen in high-income settings: as illustrated in Figure 1, this would result in a much faster or slower exponential rise in cases, with epidemic peaks occurring as soon as 3 months and as late as 7-8 months given the range of R_0 values used in the model. Furthermore, the age spectrum of disease might be considerably different in the country, for example because relevant co-morbidities are relatively more concentrated in younger age groups, as might be the case for HIV and tuberculosis. Lastly, the adjustment to CFR (mortality per case) made above is likely to not quite capture how lethal the virus would be if availability of treatment is only partial or very scarce: this is currently unknown, but it is plausible that many of the younger patients in need of respiratory support would die without oxygen support.

  Likely direction of bias: All else being equal, underestimation of severe cases, hospitalisation requirements and deaths.

- The age-specific contact matrix is not known for most countries in the world. Only a few studies have been conducted to quantify physical contacts among different age groups (for example, only 4-5 such studies exist from the African continent, though more are being done). To cope with this problem, the country reports rely on so-called 'synthetic' contact matrices, the result of separate work to adapt contact matrices for European countries to other countries in the world, based on variables (e.g. household size, schooling attendance, economic activity) that correlate with the amount of contact. It is possible that these synthetic matrices don't quite capture contact patterns in each country; in particular, for some countries it is possible that contacts of old people may be underestimated. This would not greatly affect the estimated effect of different interventions, but would result in underestimation of the overall epidemic size (total severe cases and deaths in particular).

Likely direction of bias: Underestimation of severe cases, hospitalisation requirements and deaths.

Notes on specific tables and figures

Unmitigated epidemic

Figure 1 shows the median (i.e. most likely) epidemic curve in red, for both incident (i.e. new) symptomatic cases and deaths. Each grey curve is the output of a single model run, showing the amount of variation in predictions that combinations of the statistical uncertainty distributions used for different parameters actually result in. The main parameter driving the uncertainty in these curves is the R_0: higher values of R_0 result in a steeper, earlier peak, and vice versa.

Figure 2 shows the forecast epidemic curves for different quantities. It is not surprising that symptomatic cases occur overwhelmingly among younger age groups (since these are the majority of the population in the countries currently covered by the reports), while cases requiring hospitalisation and deaths occur mainly among the elderly (since COVID-19 severity increases exponentially with age).

Table 3 will typically show cumulative outcomes plateauing by 6 months, with little further increase beyond that time point: this is because by month 6, the model predicts that an unmitigated epidemic would have mostly ‘burned out’ on its own, i.e. ended naturally once the number of people susceptible declines to such an extent (the so-called ‘herd immunity threshold’) that the epidemic goes to extinction. This is easier to visualise in Figures 4a and 4b, which both show the unmitigated curve as the leftmost on each graph.

Intervention scenarios

Table 4 should be self-explanatory: the percent reductions are relative to what would happen in the absence of any intervention. The table is somewhat redundant, in that cases needing hospitalisation and deaths are affected by the interventions in the same way. Distancing has a relatively greater effect on transmission overall than shielding, which in turn means a greater reduction in symptomatic cases; however, shielding is more impactful at reducing severe outcomes, as it specifically targets people at high risk.

Figures 3a and 3b illustrates not just the dampening of transmission that is achieved under different intervention scenarios (so-called ‘flattening the curve’), but also the extent to which interventions result in a slower evolution of the epidemics, i.e. a longer period over which a country would have to deal with some level of SARS-CoV-2 transmission. This is one of many trade-offs inherent in this pandemic.

Note that Table 6 does not show the numbers of outcomes under different scenarios, but rather the numbers averted by each intervention bundle: it’s a way to quantify the health gains, at least in terms of direct COVID-19 burden, of different response options.

Frequently asked questions

Why do the forecast deaths not include mortality indirectly due to the COVID-19 epidemic, e.g. due to disrupted livelihoods or health services?

This would require a completely different and separate model, as well as many more assumptions about what might happen to livelihoods and health services as a result of the epidemic, with each service area (e.g. maternal health, management of acute malnutrition, HIV care) requiring specific assumptions about drops in utilisation and quality, and the impact these drops would have on health outcomes. In short, a much more complex and country-specific analysis would be required. Some groups are starting to predict these indirect effects, either for specific countries or for regions of the world.

Why do the reports only cover 12 months? What might happen after 12 months?

They cover 12 months as it is difficult to presume what might happen beyond this timespan (e.g. whether game-changing vaccines or antiviral treatments might become widely available). In separate research (see https://cmmid.github.io/topics/covid19/covid-response-strategies-africa.html), we have estimated that very aggressive control interventions (e.g. periods of lockdown followed by extensive and sustained general physical distancing) could result in a 'second wave' of the epidemic after 12 months. Note that, under a lockdown scenario, this 'second wave' would come even sooner (i.e. as soon as lockdown measures are lifted) unless mitigation interventions are implemented.

Won't the epidemic look different in different parts of the country?

Yes, very probably. The forecasts in these reports don't capture this nuance. It is possible that countries might see an early, steep rise in cases in the crowded urban centres, with transmission being much slower and protracted in rural parts of the country. Specific populations, e.g. displaced people in camps, might also experience very dramatic epidemics.

Why did you not consider testing and contact tracing as one of the possible response interventions?

We wanted to emphasise interventions that low- and middle-income countries might feasibly implement. Also, testing and contact tracing is relatively harder to model and its impact depends considerably on what happens early on in the epidemic curve, when chance and small numbers play a heavy role. This intervention becomes relatively marginal as transmission intensifies.

What about the coverage and effectiveness of case management?

We have made no assumptions on these parameters. The actual capacity for hospitalisation (critical and non-critical) is not well quantified for many countries, and it isn't clear by how much it could be scaled up effectively. However, it is likely that for most countries capacity would be surpassed unless the country manages to suppress transmission until a vaccine becomes available.

How long should be a lockdown maintained?

Lockdowns are likely to work on temporarily suppressing COVID-19 epidemics, provided they are implemented effectively. They mainly postpone the epidemic, which may be valuable as a way to prepare for a more long-term response. We estimate very roughly that a two-month lockdown might delay the epidemic by around 3 months. However, this does not mean lockdowns are an appropriate intervention everywhere, as they clearly result in very high economic and societal costs.

Why are the forecasts so much higher than reported cases or deaths in my country?

Firstly, the country reports present forecasts over a much longer period than the epidemic has been active in the country: COVID-19 epidemics start with a relatively silent phase, followed by a fairly sudden exponential increase. This large increase may have been (temporarily) staved off in your country through a combination of lockdowns and early testing, or other interventions.

Secondly, it is very important to remember that reported cases and deaths reflect only the tip of the iceberg, i.e. what is visible through testing. Depending on how much testing takes place, and where testing is available, the entire country or parts of it (e.g. more rural, outlying areas) might in fact be seeing a lot more cases and deaths than the levels reported.

Thirdly, as discussed above the model is quite naïve to in-country reality.