What does preliminary data from Mumbai's second serosurvey tell us, and does it change Mumbai's COVID story?

(Murad Banaji, 05/10/20)

Some partial results from the second round of Mumbai's serosurvey have recently been released. This analysis is based on the limited data reported in the news, and may need updating when a more complete technical report becomes available. However, it seems necessary to comment at this stage, since the results are being reported in misleading ways.

The survey, completed in late August/early September in the same three wards as the first survey (F/N, M/W and R/N), found that 45% of slum residents and 18% of non-slum residents had detectable antibodies to SARS-CoV2. Assuming these are "raw" figures, not adjusted for sampling or sensitivity/specificity of the tests, these values compare to 54% in the slums and 16% in non-slum areas during the first round of the serosurvey, carried out during late June to mid-July. (Many news reports are comparing the values from the second survey to adjusted values from the first, but this could be an error.)

The mid-point of the first survey in the slum wards (based on data here) was around July 6, and in the non-slum areas was around July 13. For simplicity, the mid-point of the survey is taken to be July 9. The second survey was concluded on September 1 according to this report, and we assume a mid-point of August 23, 1.5 months after the mid-point of the first survey.

The new serosurvey has been widely (mis)reported as showing "falling infections", and as "good news". How could it be that slum prevalence fell significantly and non-slum prevalence barely rose in the space of 6 weeks? Does this data mean that earlier accounts of Mumbai's epidemic, including my own accounts here, here and here, have been built on flawed data and the story now changes? For example, if earlier estimates of prevalence were too high, then we would have to revise estimates of infection fatality rate (IFR) upwards – this would not be good news.

The short answer is, no, Mumbai's story does not appear to change dramatically in the light of the new data, and previous analyses of prevalence, fatality, and uneven spread remain largely valid. The new serosurvey data is, however, a valuable addition, and to understand how requires digging a little deeper into antibody measurements.

Possible explanations for the apparent drop in prevalence in the slums

First, for completeness, let's look at all possible explanations for the results of the second serosurvey as reported so far. Let's leave aside the non-slum areas for the moment, and consider explanations for the apparent drop in slum prevalence. Three alternatives come to mind:

1) **It is real**. There was a genuine fall in the proportion of slum-dwellers who had had COVID-19 in the surveyed wards between the two serosurveys. This would have to be via birth, death or migration.
2) **It is a consequence of sampling errors**. The sample populations surveyed in one or both serosurveys were *not* representative of the slum population of the wards. Hence sample seroprevalence did not well approximate population seroprevalence in one or both surveys. This would also suggest that the survey authors' overestimated reliability of the numbers from the first survey.
3) **It is about the tests**. The antibody test used in the serosurveys increasingly fails to pick up past infections, particularly of individuals who were infected a considerable time before the

survey. This results in falling accuracy in prevalence estimates over the duration of a long epidemic, and we see the consequences particularly starkly in measured seroprevalence from Mumbai's slums.

Of course, the reality could involve a combination of these three explanations. So we'll consider them all, even though the evidence points very strongly towards a dominant role for the last of them.

A real drop in prevalence

This could be caused by an influx of COVID-negative individuals or an efflux of COVID-positive individuals or a combination of both. At the last census, the surveyed wards had a total population of ~1.37M, and the changes required to cause a drop in prevalence of 9-12% would involve lakhs of people. Moreover, if the prevalence value of 57% from the first serosurvey is taken at face value, then with infection rife in the city between the surveys, and cases continuing to come in at an average of more than 1000 per day, we would expect slum prevalence to have risen at least a few percentage points by the time of the second serosurvey – even if the majority of spread was occurring in non-slum areas. This makes the demographic changes required to cause such a drop even more startling.

Since births or deaths could not cause changes of the scale required, we have to consider migration. There are no reports of net migration out of the slums in July-August, and in any case there are no plausible reasons why such an exodus would disproportionately involve those who have had COVID, thus lowering prevalence.

On the other hand, there has been some return of migrants to the city. Two conditions would have to hold for this to cause a drop in prevalence of the scale required: 1) the influx to these wards would have to be large – at least 2 lakhs; and 2) these migrants would have to have largely not been exposed to COVID. Even if the population of these wards increased by almost 20% in 6 weeks (which seems unlikely), prevalence among returning migrants would, one assumes, mirror prevalence in the slums during May when the great migrant exodus occurred. All the evidence is that COVID was widespread in the slums in May, and we would expect a significant proportion of returning migrants to be seropositive.

Thus, although there may be genuine fluctuations in prevalence due to inflows and outflows of people, these would likely be small, and a significant drop in prevalence even as the epidemic rages is *highly* implausible.

Sampling errors

A total of 4234 samples taken from slum dwellers were analysed in the first serosurvey, and indications are that this number was somewhat smaller in the second survey. According to one report, 5384 participants were surveyed in total across both slums and non-slum areas, while according to another report, 5840 individuals were surveyed, with 3,700 from the slums. The sampling methodology was the same in both surveys, which were designed with some care to avoid overrepresenting any particular cluster of infections in a local area. The 95% CI on the raw prevalence from the slums in the first survey was given as 52.7% to 55.6% by the authors. Suffice it to say that obtaining a raw sample prevalence of 45% from these areas at the time of the first serosurvey would be exceedingly unlikely to occur by chance.

With the passage of time and daily cases in the period between the serosurveys averaging over 1000, we would in fact expect prevalence to rise significantly, making a measurement of 45% in the second serosurvey simply on account of fluctuations in sampling even more unlikely. So, while raw sample prevalence could well under/overestimate population prevalence, it is highly implausible that this accounts for the bulk of the difference between the raw seroprevalence value in the second serosurvey and what we might expect based on the first. At most, it could play a minor part.

Waning antibodies and decreasing sensitivity of the tests

This brings us to the third explanation, that the tests are increasingly failing to pick up past infections. This last explanation in fact accords both qualitatively and quantitatively with what we know.

It is well known that antibodies to SARS-CoV2 wane quite rapidly over time, although this does not necessarily imply rapidly waning immunity – the relationship between antibody levels and protection from infection is complicated. Moreover, different test-kits are more or less sensitive to this diminishing in antibodies. A certain proportion of individuals who are measured to be seropositive with a given test-kit at some point after infection, become seronegative (they "serorevert") as far as this kit is concerned some time later. This waning of sensitivity is hard to quantify absolutely, partly because initial antibody responses vary between individuals and so the time to "seroreversion" will also vary. For example, the initial antibody response is known to be weaker in individuals with mild or asymptomatic disease. It is also possible that some individuals do not serorevert, but we do not accurately know what proportion.

When it comes to the test-kit used in the Mumbai serosurveys, the "Abbott Diagnostics Architect N-protein based test" (henceforth abbreviated to just "Abbott"), this kit has been shown in longitudinal studies to identify a fairly rapidly dwindling proportion of people who have had COVID-19 over time. Studies from New York and Brazil have attempted to quantify the waning sensitivity of this test and to develop methodologies to correct for this in order to arrive at reliable prevalence estimates based on longitudinal data obtained using the Abbott test (here and here). The second study, entitled *COVID-19 herd immunity in the Brazilian Amazon* (Buss *et al*) develops an approach to correction for waning sensitivity of the test which is applied to serosurvey data from Manaus and Sao Paolo to estimate the dynamics of the disease.

Applying the approach in Buss *et al* to Mumbai's data

The approach described in *COVID-19 herd immunity in the Brazilian Amazon* can be applied to Mumbai's data. This approach has the advantage that it is data-driven, and so it's conclusions are to some extent independent of those from more model- or extrapolation-based approaches.

The question we wish to answer, given repeated serosurvey data from a population, is: what would seroprevalence at each time point be if initial sensitivity of the test was 100% and, crucially, this did not diminish over time? We'll call this hypothetical value the "true" seroprevalence at any time point. A note of caution: the true seroprevalence as interpreted here includes individuals who may have no antibody response to infection at all (or are "serosilent"). It thus provides an estimate for prevalence after a delay accounting for the expected time for seroconversion.

The approach of Buss *et al* to estimating true seroprevalence requires both regular measurements of seroprevalence, and estimation of two parameters associated with the test: 1) the proportion of individuals who eventually serorevert for a given test, and 2) a parameter which describes the fraction of individuals who serorevert in a given time-interval after infectio – roughly speaking, the speed of seroreversion at a population level. For the parameters associated with the Abbott test, we

can use the values estimated by Buss *et al* from Brazilian data. Since these values may be somewhat context-dependent, we can explore the effects of varying these parameters on prevalence estimates.

One issue with Mumbai's serosurvey data is that data is available from only two serosurveys. To augment this data-set we imagine a third serosurvey, carried out at the end of May; the values of measured seroprevalence that *might* have been obtained in this hypothetical survey become additional parameters. Higher values of this hypothetical measured seroprevalence would correspond to an earlier epidemic, and lower values to a later epidemic. We can also have a guess at these values based on case numbers.

The details and outcomes of this approach are given in a technical appendix. We arrive at estimates for true seroprevalence in the slums/non-slum areas at the time of the second survey ranging from 62%/23% (late epidemic, smaller proportion of people serorevert, relatively slow antibody waning) to 91%/33% (early epidemic, larger proportion of people serorevert, relatively fast antibody waning).

Point estimates from the Brazilian survey and best guesses for seroprevalence at the end of May based on case data gives point estimates for "true" seroprevalence in the slums/non-slum areas of 78%/28% by the time of the second serosurvey. Both of these values would be expected to be somewhat higher today, particularly the latter, as the September surge seems to have been primarily in the non-slum areas.

These estimates accord well with prevalence estimates using case-data and the first serosurvey data. Those estimates took the first serosurvey data at face value and combined it with ward-wise data on cases to estimate infection detection in slums and non-slum areas. Plausible assumptions about changes in detection since the first serosurvey were then used to arrive at a range of prevalence estimates for late September. The two approaches both use data from the first serosurvey; but after this they diverge, with one using case data, and the other using data from the second serosurvey corrected for waning test sensitivity. It is reassuring that the two rather different approaches give estimates which are in the same ball-park (and none which are very close to the raw values from the second serosurvey!)

Fatality rates

I had shown in earlier work that Mumbai's data on COVID-19 fatalities and excess deaths, combined with data on prevalence from the first serosurvey paint a picture broadly consistent with what we know of age-dependent COVID-19 infection fatality rate (IFR) worldwide. Basically Mumbai's IFR estimated from age-dependent IFR values and Mumbai's age pyramid would lie somewhere close to 0.25% if prevalence was equal across age-groups. Greater prevalence in younger age groups, along with levels of undercounting consistent with Mumbai's excess mortality data, brought the naive IFR (i.e., IFR values ignoring undercounting) down to about half of this value.

If we were to take the prevalence estimates from Mumbai's second serosurvey at face value we would end up with naive IFR estimates – and hence true IFR estimates – considerably higher than the earlier ones. This would mean that Mumbai's true IFR was quite a lot higher than expected from European age-stratified data, and it would also make Mumbai's naive IFR an outlier amongst values inferred for Indian cities which have had serosurveys so far.

Arguments based on IFR can thus be taken as some corroborating evidence that the new seroprevalence values underestimate true prevalence.

Conclusions

Although it is possible that demographic changes and issues with sampling might play some part in explaining the apparent drop in Mumbai's prevalence between serosurveys, it is likely that the dominant effect involves antibody waning with consequent decreases in sensitivity of the Abbott test used. The observed data is consistent with a fairly early surge in the slums and decreasing sensitivity of the Abbott test as quantified in Brazilian study.

Correcting for waning sensitivity, and using best-guess values for the parameters involved, we arrive at values of around 78% and 28% for "true" seroprevalence in the slums/non-slum areas of Mumbai by the time of the second serosurvey (late August). This equates to about 54% seroprevalence city-wide by this time. If we assume that test sensitivity wanes somewhat less rapidly than inferred from the Brazilian data, and/or that fewer people serorevert over time, and/or that epidemics were less rapid, then these values come down, but the lowest plausible estimates of true seroprevalence by the time of the second serosurvey are 62% and 23% in the slums and non-slum areas respectively. This would equate to about 43% prevalence city-wide. At the upper end of the estimates, we would have true seroprevalence of 92% and 33% in the slums and non-slum areas respectively, amounting to 64% city-wide seroprevalence by the time of the second serosurvey. The slum-estimate in particular seems too high and with such high prevalence we would struggle to explain why there are still a large number of containment zones – and hence presumably active cases – in slums across the city in early October over a month later.

Acknowledgement. Many thanks to @CovidSerology for pointing me to the studies from New York and Brazil.

Technical appendix

Here we present a very abridged version of the calculations presented in Buss *et al*. Details are given in the preprint, but an abridged description here:

Time is divided into steps, taken in Buss *et al* to be months, and here to be 1.5 month intervals. The key quantities of interest are:
$\rho[n]$ = measured seroprevalence at time-step n
$r[n]$ = number of new recoveries per capita at time step n

Buss et all present the natural formula:
$\rho[n]-\rho[n-1] = r[n] - \Sigma_{k=1}^{n-1} r[k]\alpha^{n-k}p(1-\alpha)/\alpha$
which says, roughly, that new per capita seroprevalence at time-step n is a sum of new recoveries and seroprevalence lost to waning. The basic idea is, given a measured seroprevalence vector $\rho[n]$, to invert the above formula to get the vector true Δ-prevalence vector $r[n]$.

The formula uses two parameters, α, the attenuation per time step, and p, the proportion of individuals who can serorevert. The following values (95% CI) are estimated in Buss *et al*:

- α: 0.7352 (0.3236, 0.7744) is given for the *monthly* attenuation. Since we are using time-steps of 1.5 months, we raise these values to the power 1.5. Moreover, we only explore monthly values in the range (0.65, 0.7744) to avoid implausibly high seroprevalence estimates in the slums.

- p: 0.9606 (0.5784, 0.9900)

We follow Buss *et al* and use seroprevalence values adjusted for assumed sensitivity and specificity of the test rather than the raw values as our "measured seroprevalence" to input into the formula. We use three time-points spaced 1.5 months apart corresponding to:
- 1. late May (this is unknown and treated as an additional parameter)
- 2. early July (first serosurvey). Adjusted seroprevalence in slums/non-slums: 58%/17% (correction assuming sensitivity of 90% and specificity of 100%)
- 3. late August (second serosurvey). Estimated adjusted seroprevalence in slums/non-slums: 50%/20% (correction assuming test sensitivity of 90% and specificity of 100%). These values could change after sampling corrections.

We use the following values for hypothetical measured seroprevalence in late May in the slums: 10% (low), 20% (medium), 30% (high); and the non-slum areas: 2% (low), 6% (medium), 10% (high). Low values correspond to a late-developing epidemic, whereas higher values correspond to an early epidemic. Note that a seroprevalence of x% in the slums in late May would correspond to prevalence of x% in early May assuming 2-3 weeks for seroconversion on average.

Note that there had been ~30K recorded cases by May 24, as compared to ~88K by July 8. If seroprevalence was approximately tracking cases, then, in the light of the first serosurvey, we would expect about 20% seroprevalence in the slums and 6% in housing societies by late May – so the medium estimates can be regarded as our best guesses for what a serosurvey in late May might have thrown up. Varying late May's estimated seroprevalence, and varying α and p gives the following table of estimates of prevalence in slums and non-slum areas. All percentages are rounded to the nearest whole number.

		α = 0.65			α = 0.7352			α = 0.7744		
		low	medium	high	low	medium	high	low	medium	high
p=0.5784	slums	68%	70%	73%	64%	66%	68%	62%	64%	65%
	non-slums	25%	26%	27%	24%	25%	25%	23%	24%	25%
p=0.9606	slums	81%	85%	90%	74%	78%	81%	71%	74%	77%
	non-slums	29%	30%	32%	27%	28%	30%	26%	27%	28%
p=0.99	slums	82%	87%	92%	75%	79%	82%	71%	75%	78%
	non-slums	29%	31%	33%	27%	28%	30%	26%	27%	29%

Highest estimates highlighted in red, lowest in green, and values based on best-guesses on each parameter are in yellow. As 78%/28% seem somewhat high for slum/non-slum prevalence in late August, it is possible that Mumbai's monthly attentuation is somewhat lower than estimated from Brazilian data and/or the proportion of individuals who serorevert is lower.

From the calculations we also obtain estimates of the true seroprevalence at the time of the first serosurvey. The best-guess parameters correspond to slum/non-slum true seroprevalence of 65%/19% at the time of the first serosurvey: compare these to the adjusted values of 58%/17%.