 doi:10.20944/preprints202012.0126.v1

Understanding the challenges and uncertainties of seroprevalence studies for SARS-CoV-2

David McConnell* [a,b], Conor Hickey [a,b], Norma Bargary [c], Lea Trela-Larsen [a,b], Cathal Walsh [a,c], Michael Barry [a,b] Roisin Adams [a,b]

*Corresponding author, DMcConnell@stjames.ie

[a] National Centre for Pharmacoeconomics, St James's Hospital, Dublin, Ireland

[b] Department of Pharmacology and Therapeutics, Trinity College Dublin, Dublin, Ireland

[c] Health Research Institute and MACSI, University of Limerick, Limerick, Ireland

Abstract

SARS-CoV-2 continues to widely circulate in populations globally. Underdetection is acknowledged and is problematic when attempting to capture the true prevalence. Seroprevalence studies, where blood samples from a population sample are tested for SARS-CoV-2 antibodies that react to the SARS-CoV-2 virus, are a common method for estimating the proportion of people previously infected with the virus in a given population. However, obtaining reliable estimates from seroprevalence studies is challenging for a number of reasons, and the uncertainty in the results is often overlooked by scientists, policy makers and the media. This paper reviews the methodological issues that arise in designing these studies, and the main sources of uncertainty that affect the results. We discuss the choice of study population, recruitment of subjects, uncertainty surrounding the accuracy of antibody tests themselves, and the relationship between antibodies and infection over time. Understanding these issues can help the reader to interpret and critically evaluate the results of seroprevalence studies.

Keywords: Covid-19; SARS-CoV-2; coronavirus; seroprevalence; antibody testing

Introduction

The SARS-CoV-2 virus is likely to be circulating in populations since December 2019 with the currently known first case recorded in Wuhan in China. Despite worldwide attempts at suppression, and in some countries eradication, the virus continues to circulate and in many countries it is unclear to what extent. However many publications have indicated that it is circulating to a wider extent than the case incidence report[1–3].

Estimating the true number of people who have previously been infected with SARS-CoV-2 will enable scientists and policy-makers to understand how the virus spreads in various settings; to estimate the proportion of people infected who will need hospital care and/or ICU admission, and thus plan for further outbreaks; and to evaluate the effectiveness of restrictions aimed at curbing the spread of the virus. The quantification of those who have been exposed to infection is ideally done via direct testing for presence of the virus. However, this is dependent on capturing those

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

cases while a person is shedding virus. An alternative method is estimating the exposure of the virus via the presence of antibodies.

One approach to doing this is to carry out 'seroprevalence' (or 'antibody') studies, which are used to estimate the proportion of people who have SARS-CoV-2 antibodies in their blood. Seroprevalence studies have been the subject of much recent attention, however estimates of total case numbers arising from these studies can be highly uncertain. Overlooking the uncertainties and limitations of these studies can lead to a flawed understanding of the disease and its spread, and subsequently, poor policy decisions. The purpose of this article is to give an informal explanation of some of the main sources of uncertainty and areas of misunderstanding that can arise in such studies.

The problem with confirmed cases

As of the 27th November 2020, there were a total of 71,493 confirmed cases of COVID-19 in Ireland[4]. There is good reason to believe that the true number of those who have been infected is considerably higher than this, as outlined in Figure 1. Many of those who were infected were not tested at all, particularly in the earlier stages of the epidemic in Ireland, when testing was largely restricted to those showing two or more 'typical' symptoms, healthcare workers, and high-risk groups. While testing is more widespread now, it is still likely that some mild or asymptomatic infections remain undiagnosed. Additionally, the tests used to detect active infection are not perfect[5] – according to one study[6], a currently infected person has at most a 67% chance of correctly testing positive (this figure can be much lower, depending on the length of time since exposure). This means that many infected people are recorded as testing negative and are not counted in the confirmed cases figures. While false positives (uninfected people testing positive) can also occur, these are likely to have a negligible impact by comparison. As a result, confirmed cases almost certainly underestimate total infection numbers to varying degrees around the world.

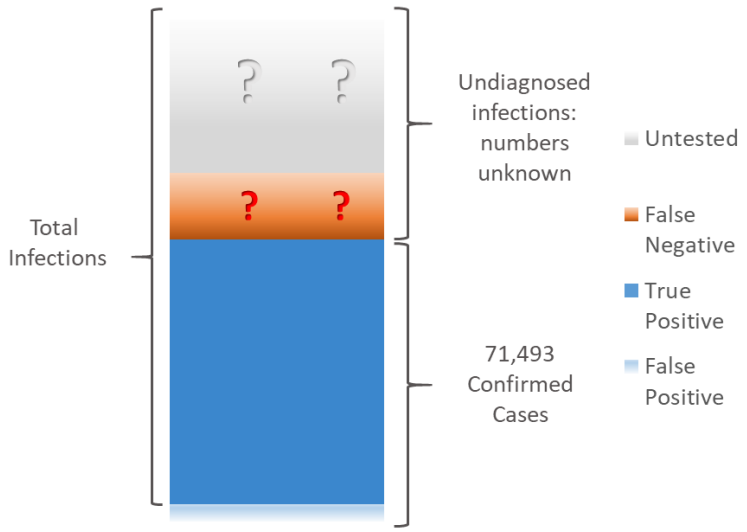

Figure 1: Total COVID-19 Cases in Republic of Ireland as of 27th November 2020. Official COVID-19 case numbers (HPSC[4] 27th November 2020) are based on positive tests only (a small number of which may be false positives). The true number infections also includes an unknown number of people who were either never tested, or else falsely tested negative.

What is a seroprevalence study?

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

When infected with SARS-CoV-2, a person's immune system produces antibodies to fight the virus, which, in the majority of cases, become detectable typically within 14 to 21 days of infection[7] and remain detectable for at least several months afterwards[8,9]. There are a number of tests available to detect the presence of SARS-CoV-2 antibodies in a recovered patient's blood. By carrying out antibody tests on a group of people, we can determine who has been previously infected. These tests enable us to carry out studies to estimate seroprevalence in a population, that is, the proportion of people with SARS-CoV-2 antibodies in their blood, and thus estimate the total number of previous infections.

Many antibody studies have been carried out around the world. While the results have varied considerably between studies and locations, they have consistently indicated that the true number of people previously infected is considerably higher than the official number of confirmed cases[10]. For example, since June 2020 the HSE's Health Protection Surveillance Centre (HSPC) has been carrying out a seroprevalence study (SCOPI) in Ireland[11]. Preliminary results published in August 2020 estimated that 1.7% of the population aged 12-69 years in Ireland had previously been infected, corresponding to 59,500 total infections, which was approximately 3 times higher than the total number of confirmed cases in this age bracket at the time.

Unless we test everybody in the population, using a perfect (i.e. 100% accurate) test, we cannot calculate seroprevalence exactly: we can only estimate it alongside the associated uncertainty. Some of this uncertainty can be quantified, and this is usually presented (in the form of 'confidence intervals,' which we discuss later) in the results. For example, the SCOPI study suggested that a range of 1.1% to 2.4% seroprevalence in Ireland was plausible, corresponding to between 39,800 and 85,200 infections in this age group. However, confidence intervals like this capture some but not all of the uncertainty in the study results, and it can therefore be challenging for the non-expert to make sense of these figures. In what follows, we discuss some of the major sources of this uncertainty: issues around whom we include and exclude from the study, how we recruit participants, the role of random chance, and the limitations of antibody tests themselves. Understanding these issues can help the reader to avoid drawing misleading conclusions.

Populations and generalization – Seroprevalence among whom?

In a seroprevalence study, the group of people among whom we are trying to estimate prevalence of specific antibodies is called the *population*. Population could mean many things in this context: the Irish population as a whole, healthcare workers in Leinster, nursing home residents in Cork, etc. One of the easiest ways to draw flawed conclusions from a study is by incorrectly generalizing the results beyond the group of interest. The population is typically much larger than the group of people that we actually test.

For example, the SCOPI study enrolled people aged 12-69 [11]. It may be tempting to conclude that the prevalence of past SARS-CoV-2 infection in the over 70s will be similar to that in the study population, however, there are a number of reasons why this may not be the case. For example, the government's cocooning advice may have reduced infection rates in this age group, while on the other hand, it is likely that hospital and nursing home outbreaks disproportionately affected the over 70s. Similarly, seroprevalence studies carried out in England [12], Scotland[13], the Netherlands[14] among many others use samples collected from blood donors. Many people with long term medical conditions are not eligible to donate blood and thus would not be included. Moreover, even among those who are eligible to donate blood, those who actually do donate may differ in important ways from those who do not - perhaps healthcare workers are more likely to donate blood, or maybe blood donors typically follow social distancing guidance more than the average person. It is not

 doi:10.20944/preprints202012.0126.v1

immediately obvious whether or not donors will be more or less likely to have been exposed than the general population.

The extent to which this matters depends on how we intend to use the results: for example, given that the majority of COVID-19 deaths in Ireland occurred in the over 70s age group, it would be unwise to draw any firm conclusions about the overall infection fatality ratio (IFR) in Ireland from the SCOPI study (though it could be used to estimate the IFR among the under 70s). In general, it is essential for the reader to be aware of the population in which a seroprevalence has been carried out, as this is often more restrictive than we would like. While it may be case that seroprevalence is indeed similar in groups outside the study population (e.g. in the over 70s or among non-blood donors respectively, in the two examples just discussed), drawing such a conclusion requires making assumptions, which represent an additional source of uncertainty that cannot be easily measured. The plausibility of these assumptions should be assessed using qualitative judgement and external information if possible.

Selecting good samples

Typically, we cannot test everybody in a given population (except in cases where the population is small and easily identified, such as the residents of a specific nursing home). Instead, we select a *sample* of people to test, and based on the results, try to infer the prevalence in the population as a whole. However, we can only do this in a meaningful way if the sample is representative of the wider population, in terms of characteristics that affect the likelihood of having been previously infected.

As an example, imagine that we selected a sample of the population that was predominantly female (Sample 1 in Figure 2). While there may be no biological reason to suggest that either men or women are more or less likely to be infected (although there is some evidence that men are more likely to experience severe disease), the prevalence of SARS-CoV-2 in a sample like this could differ from the general population for any number of reasons. For instance, it is quite plausible that workers in sectors that are predominantly female (e.g. healthcare) face a higher risk of being infected than the workforce at large. If these workers are over-represented in the sample, then the prevalence of SARS-CoV-2 may be over-estimated by this study. By contrast, in a sample with approximately equal numbers of male and female subjects (Sample 2 in Figure 2), this issue does not arise, and seroprevalence in the sample is much more likely to be close to the true (population) value.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

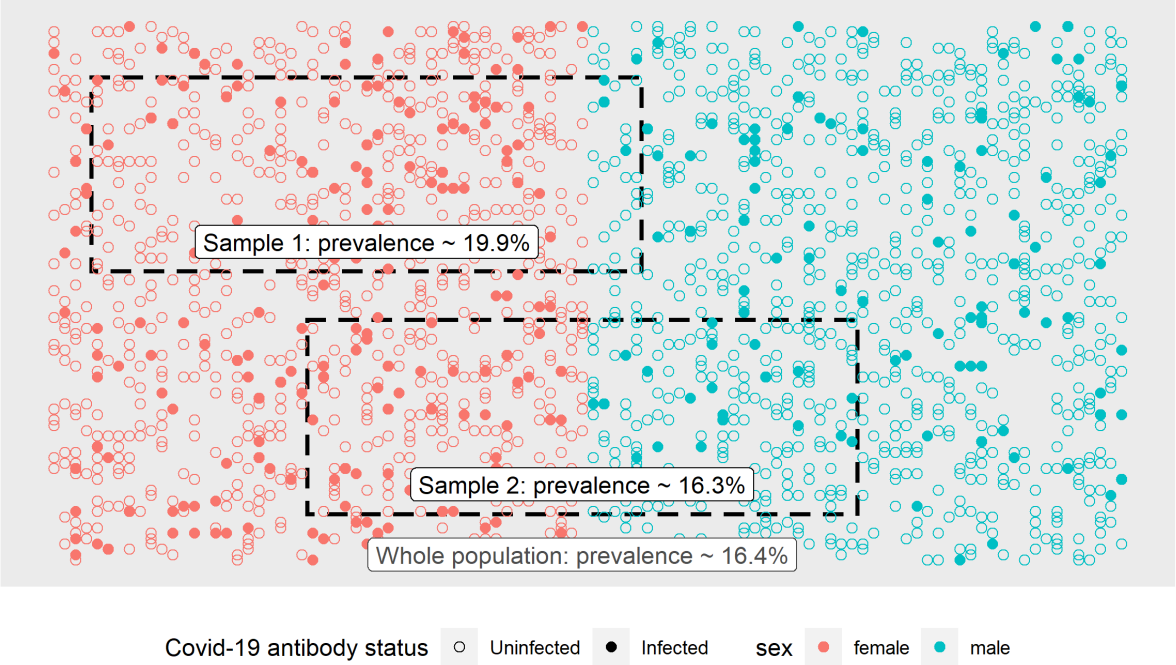

Figure 2: Importance of representative sampling. In this population, the rate of prior infection is higher among females. As a result, Sample 1, which is predominantly female, overestimates the prevalence of prior infection among the wider population. By contrast the prevalence in Sample 2, which contains approximately equal numbers of male and female subjects, is much closer to that of the wider population.

There are many less obvious ways in which sample selection can introduce bias into prevalence estimates. For example an early antibody study in Santa Clara, California[15] recruited participants using targeted Facebook ads. This could be a source of bias if Facebook users differ systematically from the population as a whole, and in ways that affect the likelihood of COVID-19 exposure; age profile for example. Other studies[16] have enrolled participants outside shopping centres. This approach can be better, provided that these areas cover a broad spectrum of the population in terms of age, sex, socioeconomic status and other factors. When studies are carried out with the involvement of government agencies, it is often possible to select a representative sample of the population from an official 'list' such as the electoral register or similar, and inviting those selected to participate. For example, the REACT study in England[17] selected participants at random from the NHS patient list, which contains everybody registered with a GP in England. Having an (almost) complete list like this from which to choose a sample is ideal, as it is much less likely to systematically exclude any large groups of the population, though this option is not always available to the study investigators.

Irrespective of who we invite to participate in a study, it will not necessarily be the case that all of those selected will agree to participate. This raises an important question – do the infection rates differ among those who agreed to participate and those who did not? For example, in some studies[15] participants are informed of their antibody test results afterwards: it is possible that those who have recently experienced flu-like symptoms would be more likely to seek a test, compared with others in the population who had not. More generally, agreement to participate in a study could be influenced by a number of factors that in turn are related to the likelihood of previous exposure: socioeconomic status, trust in government or the medical profession, and many others.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

Some of these challenges in sample recruitment can be accounted for in the study design, however, many cannot. In this case, the usual approach is to carry out an adjustment on the results, in order to reflect what would be obtained from a sample that 'looks like' the population of interest in terms of pre-determined characteristics (age, employment status, employment sector, living arrangements and so on). There are a number of methods for doing this, and many well-designed studies will carry out some form of adjustment. However, this is not easy to do well: the reliability of the results depends on choosing appropriate factors to adjust for.

Ultimately, sample selection is a key challenge that affects nearly all seroprevalence studies – if the sample differs from the wider population in ways that affect the likelihood of exposure to the virus, then the estimated seroprevalence may be unreliable.

Uncertainty from sampling

No matter how carefully we choose our sample, it is unlikely that the prevalence of past infection among those sampled will be exactly the same as that of the wider population – some variation is inevitable due to random chance. Imagine that the true prevalence of past SARS-CoV-2 infection in the population is 5%, and that we randomly selected a sample of 100 people from the population. We would not be hugely surprised if the number of people in the sample previously infected were not exactly 5. It is easy to imagine such a sample containing 1, 3, 10, maybe even 12, simply due to chance. Thus in a seroprevalence study of only 100 people, we could quite easily obtain very inaccurate estimates of say 1% or 12%, from which we could draw misleading conclusions.

In a much larger sample of say 100,000 people, we still wouldn't expect exactly 5,000 (5%) of them to have been previously infected: however, finding as few as 1,000 or as many as 12,000 would be extremely unlikely. As sample size increases, estimates that are very wrong due to chance become less likely, and thus we can be more confident that our results are close to the true value - provided the sample has been selected appropriately. In fact, based on the sample size, we can measure uncertainty due to random variation: we can calculate how far away the estimate (i.e. the prevalence in the sample) could reasonably be from the true value (the prevalence in the wider population), subject to an acceptable level of uncertainty.

To describe this, results of seroprevalence studies consist of two elements:

- a *point-estimate,* i.e. a single 'central' or 'most-likely' value for the percentage of people who are estimated to have been previously infected.
- an *interval-estimate,* or a range of values surrounding the point estimate. This could be called a confidence/credible/uncertainty interval (the differences between these are quite technical and will not be discussed here). They represent the range in which we expect the 'true' value of prevalence to lie, with a reasonable degree of certainty, and we consider values outside this interval to be unlikely. Thus a wider interval suggests a greater degree of uncertainty compared with a narrow one.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

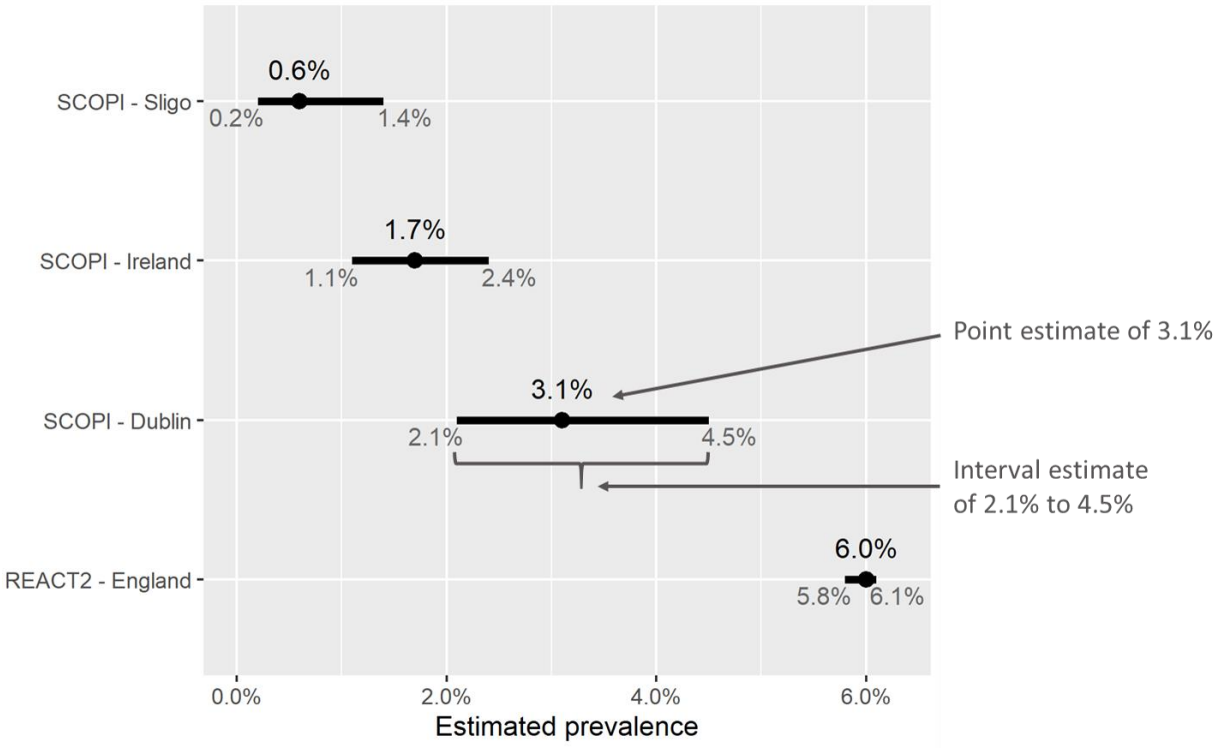

Figure 3: Point- and interval estimates of seroprevalence from two studies, SCOPI[11] in Ireland and REACT2[18] in England. Based on a sample of 913 people in Dublin the SCOPI study estimated seroprevalence at 3.1%. The corresponding interval estimate indicates that the true value of seroprevalence in the wider population in Dublin is likely to lie between 2.1% and 4.5%. By contrast the REACT2 study in England enrolled a large number of participants (ca. 100k), and thus the corresponding confidence interval for seroprevalence is narrow - greater sample sizes give more precise estimates (provided they are indeed representative of the population).

Figure 3 shows some examples of point and interval estimates of prevalence obtained from two different studies, SCOPI[11] in Ireland and REACT2[18] in England. Readers will sometimes focus on the point estimate, and ignore the width of the corresponding interval estimate, and media coverage of these studies typically reflects this. This can be a mistake, particularly when intervals are wide, as other values of prevalence (other than that given by the point estimate) may be likely. For example in Figure 3, the upper end of the interval estimate for seroprevalence in Sligo (1.4%) is seven times higher than that of the lower estimate (0.2%); we can conclude that seroprevalence in Sligo was likely quite low (below 1.5%) but not zero, though we cannot be much more precise than this. It is always better to focus on the 'plausible range' of prevalence values implied by a study, rather than focusing on the point estimate alone.

Confidence intervals account for uncertainty due to random variation, such as the difference in prevalence between the sample measured and the wider population. As such, studies enrolling larger numbers of people will result in less uncertainty, i.e. narrower intervals. However, it is equally important to note that *not all sources of uncertainty will be captured in these interval estimates*. For example, if the sample of people that we have recruited is not representative of the wider population to begin with (and we do not adjust it accordingly), then the true population prevalence could still be quite far outside the interval.

The limitations of antibody tests

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

Like all medical tests, no SARS-CoV-2 antibody test is 100% accurate. There are two different ways in which we want an antibody test to be accurate: those who have been previously infected should test positive, and those who have never been infected should test negative. As such, test accuracy is usually described in terms of two separate measurements (Figure 4):

- *Sensitivity*: the proportion of previously infected people who will correctly test positive
- *Specificity*: the proportion of never-infected people who will correctly test negative.

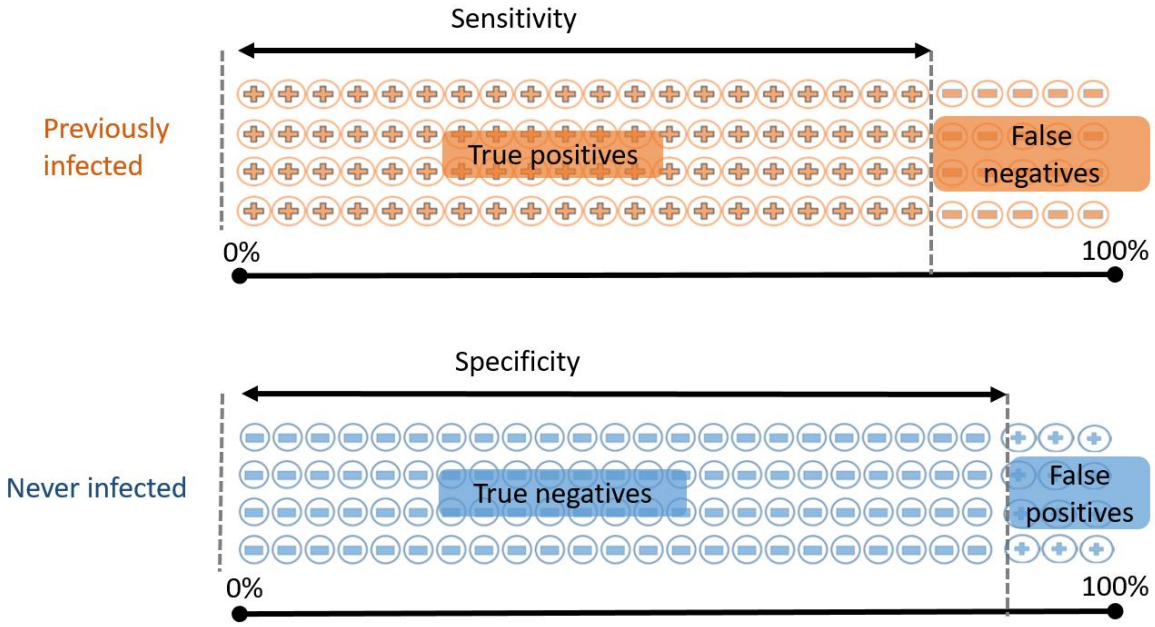

Figure 4: Sensitivity and Specificity of a Sars-Cov-2 antibody test. In a sample of previously infected people, some will correctly test positive (true positives) while others will incorrectly test negative (false negatives) – sensitivity refers to the proportion of previously infected people who correctly test positive. Similarly, never-infected people may correctly test negative (true negatives), or incorrectly test positive (false positives) – specificity is the proportion of never-infected people who correctly test negative.

Consider a hypothetical SARS-CoV-2 antibody test that has 80% sensitivity and 94% specificity. This means that

- If we test 100 people who have all previously been infected, we expect 80 of them to test positive
- If we test 100 people who have never been infected, we expect 94 of them to test negative.

To see why this matters we imagine an antibody study of 1,000 people, 50 of whom (5%) have previously been infected with SARS-CoV-2. We use this same test with 80% sensitivity and 94% specificity. The results are broken down in Figure 5. One aspect jumps out: there are more false positives than true positives, and estimated prevalence (9.7%) is considerably higher than true prevalence (5%). This occurs because:

- Of the 50 people who have been infected, 40 correctly test positive (80% sensitivity)

- Of the 950 people who have not been infected, 893 correctly test negative (94% specificity), and thus 57 test positive.

Adding these together we see that 97 people in total test positive, while the study contained 50 true cases - thus our estimate is almost twice the true value.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

Figure 5: The outcome of testing 1,000 people, 50 of whom have previously been infected, using an antibody test with 80% sensitivity and 94% specificity. The true prevalence is thus 5% (50 out of 1000), while the apparent prevalence, i.e. the proportion of tests that give a positive result, is 9.7%.

The magnitude of how far the estimated seroprevalence will be from the true prevalence depends not only on the test, but also on the true prevalence in the population. Table 1 shows the results obtained using the same test as before on a sample of 1,000 people. If the true prevalence is 2%, a study like this gives an estimated prevalence of 7.5% which is almost 4 times higher – an overestimation of this magnitude can lead us to very misleading conclusions. On the other hand, if the true prevalence is 20%, the test performs much better, with an estimate of 20.8% prevalence, which may well be 'close enough' for many purposes. These examples show that test accuracy can potentially make a big difference to the estimated prevalence numbers and lead to very misleading results in some cases. These issues should not be overlooked when studying the spread of infection and developing related policy – for example, the potential impact of imperfect test accuracy on estimated seroprevalence should be explored, particularly in cases where the results are surprising.

Table 1: Two examples of estimated seroprevalence using a test with sensitivity 80% and specificity 94%.

	Example 1: true prevalence 2%			**Example 2: true prevalence 20%**		
	Negative	**Positive**	**Total**	**Negative**	**Positive**	**Total**
Never infected	921	59	980	752	48	800
Previously Infected	4	16	20 (2% true prevalence)	40	160	200 (20% true prevalence)
Total:	925	75	1,000	792	208	1,000

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

		(7.5% estimated prevalence)			(20.8% estimated prevalence)	

Challenges when correcting for imperfect test performance

When faced with imperfect tests in seroprevalence studies, we have two options:

- Accept that this is a limitation of the study, and that the true prevalence might therefore differ from what was estimated in the study. In this case, we are estimating the proportion of the population who would test positive, using the same test.
- Try to adjust or correct the results to account for imperfect test accuracy (i.e. attempt to subtract the false positives and add the false negatives). In this case, we are estimating the proportion of the population who have SARS-CoV-2 antibodies in their blood.

For example, in a study[12] of healthy blood donors in England, unadjusted seroprevalence at the end of September 2020 was estimated at 5.7% (with a confidence interval of 5.2% to 6.3%), while adjusting for sensitivity and specificity gave an estimate of 6.1% (confidence interval 5.4% to 6.8%). The second approach seems preferable, as it aims to adjust for the problems caused by imperfect test performance and estimate the true quantity of interest – the prevalence of antibodies. Accounting for this uncertainty typically results in wider interval estimates. The main reason for this is that we do not know the 'true' values of sensitivity and specificity for the antibody test that we are using. These quantities also need to be estimated using appropriately designed studies, which again come with all the challenges of sample selection, and the uncertainty arising from random variation. On top of this, test accuracy will almost certainly vary between studies due to differences in populations (factors such as age, disease severity, time since infection can affect test accuracy), and again, simply due to chance.

Figure 6 shows estimated sensitivity and specificity, together with associated confidence intervals for three commercially available tests, taken from the U.S. Food & Drug Administration website[19]. It is worth pointing out that a recent review of studies evaluating antibody tests noted that many of them were of low quality[20], and therefore it is quite possible that the intervals in Figure 6 actually underestimate uncertainty. To account for this we have also included a fourth (hypothetical) test with lower sensitivity and specificity, for illustrative purposes.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

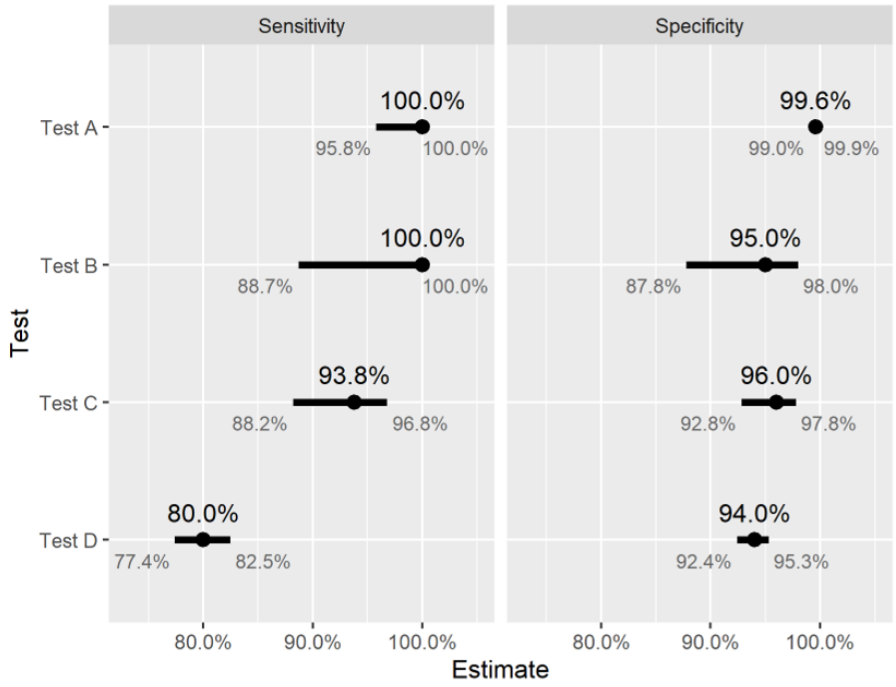

Figure 6: Estimated sensitivity and specificity of four different antibody tests. Tests A, B and C are based on data presented on the US Food and Drug Administration website (for three different commercially available tests), Test D is for illustrative purposes. Estimated sensitivity and specificity for Test A are both high, with a high degree of certainty (i.e. narrow interval estimates). For Tests B and C, estimated sensitivity and specificity are more uncertain, while for Test D these values are definitely lower.

Studies carried out using highly accurate tests (i.e. those with point-estimates of sensitivity and specificity near 100%, and narrow interval estimates for these values) will give the most precise estimates of seroprevalence in the population, compared to those using antibody tests with low or uncertain estimates of sensitivity and specificity. These factors limit how precisely we can estimate true seroprevalence in the population and can vary considerably between tests. Figure 7 illustrates how this occurs by plotting the results of 10,000 'simulated' (i.e. random, computer generated) studies carried out using different antibody tests. In each simulation, we pick a random sample of 2500 from a very large population, in which the true seroprevalence is 10%, and test them using Tests A, B, C and D from Figure 6. We then 'correct' the results to account for sensitivity and specificity and estimate the seroprevalence. The prevalence of antibodies in each sample differs randomly from that of the true population, as do estimated sensitivity and specificity, to account for what would happen in real-life.

The graphs show the distribution or 'spread' of the results from these simulated studies and illustrate how much these typically differ from the true prevalence of 10%. In each plot, the horizontal axis represents the estimated prevalence, after adjusting for test sensitivity and specificity. The height of the bars represents the number of simulated studies resulting in that particular estimate: thus if the bars are clustered close together, the corresponding test is likely to give very precise estimates of seroprevalence that are close to the true value of 10%. These estimates are centered around the true value of 10% no matter which test we use. However, those for Test A are typically much closer to the true value than any of the other tests; this reflects the fact that Test A is both more accurate (i.e. sensitivity and specificity are closer to 100%), and also that we are more certain about these values (see Figure 6). Note that with Test D, even though we have very precise estimates of sensitivity and specificity, these values are quite low: as a result, estimates of seroprevalence are quite variable.

Preprints (www.preprints.org) | NOT PEER-REVIEWED | Posted: 7 December 2020 doi:10.20944/preprints202012.0126.v1

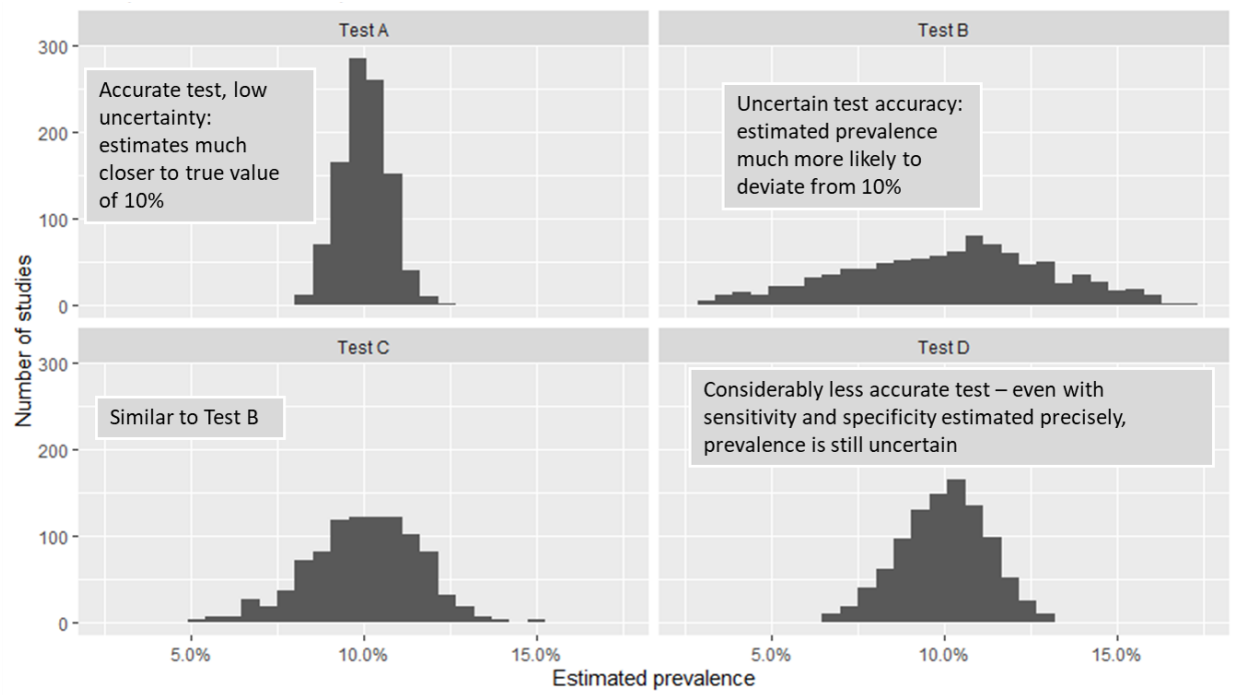

Figure 7: Estimates of seroprevalence from 10,000 simulated studies, using different antibody tests and adjusting results for imperfect sensitivity and specificity. Each simulation selects a random sample of 2500 participants from a population with an overall prevalence of 10%. The extent to which estimated prevalence differs from true prevalence on average depends on the test used - lower and/or more uncertain values of sensitivity and specificity result in more uncertain estimates of prevalence.

What all of this means is that while we can correct seroprevalence results to account for test accuracy, doing so involves incorporating additional uncertainties, which results in wider confidence intervals, i.e. less precise estimates. As a result, the precision with which we can estimate true seroprevalence is limited by the sensitivity and specificity of the antibody tests used in the study, as well as the extent of our existing knowledge about these values.

Antibodies, infection, and immunity

In many studies, the presence of specific antibodies is assumed to be a reasonable marker of previous SARS-CoV-2 infection. However, there are questions about the long-term persistence of antibodies in the blood of recovered COVID-19 patients, with some studies reporting that antibodies decrease to undetectable levels over time[21,22] (particularly among mild and asymptomatic patients), though other studies have shown the opposite[9]. This suggests that seroprevalence surveys will do a better job at identifying recent infections than those that occurred earlier in time but will inevitably 'miss' an unknown proportion of cases. Similarly, the relationship between antibodies and immunity remains unclear: it is not known whether or not the presence of SARS-CoV-2 antibodies indicates immunity, nor whether the absence of antibodies is evidence of no immunity.

This highlights a clear limitation of using seroprevalence studies as way of counting the total number of previous infections in a population: those without antibodies will not be counted. Using these studies to measure the extent of immunity within the population also requires making assumptions that we cannot yet verify. Despite these issues, seroprevalence studies are still usually the best approach to estimate previous infections and current levels of population immunity. However, it is important to remember that antibodies are only one part of the picture, and should be used together with other external information, such as confirmed cases, deaths, and infectious disease models, to better understand the disease.

Conclusion

Estimating the number of previous SARS-CoV-2 infections in a given population is important for understanding the nature of the disease and its transmission, as well as the effectiveness of strategies used to control it. While seroprevalence studies are a well-established approach to doing this, there are many challenges in estimating this number accurately, and any such estimate will be uncertain (though to what extent can vary considerably between studies). Indeed, the issues discussed in this article are not unique to seroprevalence studies, and can arise in any scientific study where we recruit a sample of people in order to estimate something in a wider population. Generally, interval estimates such as confidence intervals allow us to measure the likely effect of random variation on the results, however, they rarely capture *all* of the uncertainty that arises when trying to use these results to better understand the disease or to make decisions. For the reader, it is important to be aware of other sources of uncertainty: issues around restricted populations, non-representative samples, and importantly, what exactly is being measured and how (e.g. in our case, uncertainty around test accuracy, and limited knowledge about antibodies in the long term). Awareness of the issues discussed in this article allows us to consider these estimates with greater knowledge of how this uncertainty can arise, and to identify how it may affect our interpretation of the results.

References

1. Havers FP, Reed C, Lim T, et al. Seroprevalence of Antibodies to SARS-CoV-2 in 10 Sites in the United States, March 23-May 12, 2020. *JAMA Intern Med*. Published online July 21, 2020. doi:10.1001/jamainternmed.2020.4130

2. Russell TW, Golding N, Hellewell J, et al. Reconstructing the early global dynamics of under-ascertained COVID-19 cases and infections. *medRxiv*. Published online September 22, 2020:2020.07.07.20148460. doi:10.1101/2020.07.07.20148460

3. Flaxman S, Mishra S, Gandy A, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature*. 2020;584(7820):257-261. doi:10.1038/s41586-020-2405-7

4. Health Service Executive. Epidemiology of COVID-19 in Ireland - Health Protection Surveillance Centre. Published November 27, 2020. Accessed November 30, 2020. https://www.hpsc.ie/a-z/respiratory/coronavirus/novelcoronavirus/casesinireland/epidemiologyofcovid-19inireland/

5. Böger B, Fachi MM, Vilhena RO, Cobre AF, Tonin FS, Pontarolo R. Systematic review with meta-analysis of the accuracy of diagnostic tests for COVID-19. *Am J Infect Control*. Published online July 10, 2020. doi:10.1016/j.ajic.2020.07.011

6. Kucirka LM, Lauer SA, Laeyendecker O, Boon D, Lessler J. Variation in False-Negative Rate of Reverse Transcriptase Polymerase Chain Reaction–Based SARS-CoV-2 Tests by Time Since Exposure. *Ann Intern Med*. Published online May 13, 2020. doi:10.7326/M20-1495

7. Deeks JJ, Dinnes J, Takwoingi Y, et al. Antibody tests for identification of current and past infection with SARS-CoV-2. *Cochrane Database of Systematic Reviews*. 2020;(6). doi:10.1002/14651858.CD013652

8. Isho B, Abe KT, Zuo M, et al. Persistence of serum and saliva antibody responses to SARS-CoV-2 spike antigens in COVID-19 patients. *Science Immunology*. 2020;5(52). doi:10.1126/sciimmunol.abe5511

9. Gudbjartsson DF, Norddahl GL, Melsted P, et al. Humoral Immune Response to SARS-CoV-2 in Iceland. *N Engl J Med*. Published online September 1, 2020:NEJMoa2026116. doi:10.1056/NEJMoa2026116

10. Arora RK, Joseph A, Wyk JV, et al. SeroTracker: a global SARS-CoV-2 seroprevalence dashboard. *The Lancet Infectious Diseases*. 2020;0(0). doi:10.1016/S1473-3099(20)30631-9

11. Health Service Executive. SCOPI: COVID-19 antibody research study - HSE.ie. Published 2020. Accessed August 31, 2020. https://www.hse.ie/eng/services/news/newsfeatures/scopi-covid-19-research-project/

12. Public Health England. *Weekly Coronavirus Disease 2019 (COVID-19) Surveillance Report*.; 2020. Accessed December 1, 2020. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/923668/Weekly_COVID19_Surveillance_Report_week_40.pdf

13. Thompson CP, Grayson N, Paton R, et al. *Detection of Neutralising Antibodies to SARS Coronavirus 2 to Determine Population Exposure in Scottish Blood Donors between March and May 2020*. Infectious Diseases (except HIV/AIDS); 2020. doi:10.1101/2020.04.13.20060467

14. Slot E, Hogema BM, Reusken CBEM, et al. *Herd Immunity Is Not a Realistic Exit Strategy during a COVID-19 Outbreak*. In Review; 2020. doi:10.21203/rs.3.rs-25862/v1

15. Bendavid E, Mulaney B, Sood N, et al. *COVID-19 Antibody Seroprevalence in Santa Clara County, California*. Epidemiology; 2020. doi:10.1101/2020.04.14.20062463

16. Rosenberg ES, Tesoriero JM, Rosenthal EM, et al. Cumulative incidence and diagnosis of SARS-CoV-2 infection in New York. *Annals of Epidemiology*. 2020;48:23-29.e4. doi:10.1016/j.annepidem.2020.06.004

17. Riley S, Atchison C, Ashby D, et al. REal-time Assessment of Community Transmission (REACT) of SARS-CoV-2 virus: Study protocol. *Wellcome Open Res*. 2020;5:200. doi:10.12688/wellcomeopenres.16228.1

18. Ward H, Atchison CJ, Whitaker M, et al. Antibody prevalence for SARS-CoV-2 in England following first peak of the pandemic: REACT2 study in 100,000 adults. *medRxiv*. Published online August 21, 2020:2020.08.12.20173690. doi:10.1101/2020.08.12.20173690

19. U.S. Food & Drug Administration C for D and R. EUA Authorized Serology Test Performance. *FDA*. Published online August 17, 2020. Accessed September 2, 2020. https://www.fda.gov/medical-devices/coronavirus-disease-2019-covid-19-emergency-use-authorizations-medical-devices/eua-authorized-serology-test-performance

20. Lisboa Bastos M, Tavaziva G, Abidi SK, et al. Diagnostic accuracy of serological tests for covid-19: systematic review and meta-analysis. *BMJ*. Published online July 1, 2020:m2516. doi:10.1136/bmj.m2516

21. Long Q-X, Tang X-J, Shi Q-L, et al. Clinical and immunological assessment of asymptomatic SARS-CoV-2 infections. *Nat Med*. 2020;26(8):1200-1204. doi:10.1038/s41591-020-0965-6

22. Ibarrondo FJ, Fulcher JA, Goodman-Meza D, et al. Rapid Decay of Anti–SARS-CoV-2 Antibodies in Persons with Mild Covid-19. *N Engl J Med*. 2020;383(11):1085-1087. doi:10.1056/NEJMc2025179