The Ministry of Health's modeling of the impact of the Coronavirus on New Zealand: A look behind the headlines

April 2020

tailrisk economics

te maha hoki o nga pihikete

About Tailrisk economics

Tailrisk economics is a Wellington economics consultancy. It specialises in the economics of low probability, high impact events including financial crises and natural disasters. Tailrisk economics also provides consulting services on:

- The economics of financial regulation
- Advanced capital adequacy modelling
- Stress testing for large and small financial institutions
- Regulatory compliance for financial institutions
- General economics.

Tailrisk is prepared to undertake economics analyses of public policy proposals on a discounted or pro bono basis.

Principal Ian Harrison (B.C.A. Hons. V.U.W., Master of Public Policy SAIS Johns Hopkins) has worked with the Reserve Bank of New Zealand, the World Bank, the International Monetary Fund and the Bank for International Settlements.

Contact:
Ian Harrison – Principal Tailrisk economic
harrisonian52@gmail.com
Ph. 022 175 3669
04 384 8570

The Ministry of Health's modeling of the impact of the Coronavirus on New Zealand: A look behind the headlines

Part one: Introduction

The only publicly available information on the Ministry of Health's (the Ministry) modelling of the impact of Coronavirus on New Zealand is a set of reports on the Ministry coronavirus website, all produced by the University of Otago Covid-19 Research Group (OCRG). These papers address a range of issues and the later ones used a model that is in the public domain and is available at covidsim.eu. Here we focus on one paper -'Potential Health Impacts from the COVID-19 Pandemic for New Zealand if Eradication Fails: Report to the NZ Ministry of Health' that was dated 23 March. It has received a substantial amount of media attention.

As we understand it, the report explored the consequences if the government's actions are not successful in eliminating the virus.

The key 'headline' result in the report is that if the lockdown shock therapy fails the consequences are serious. **Between 8560 and 14400 could die over the next year.** The conclusion was:

If New Zealand fails with its current eradication strategy toward COVID-19, then health outcomes for New Zealand could be very severe. If interventions were intense enough however, in some scenarios the epidemic peak could still be suppressed or pushed out to the following year (at which time a vaccine may be available

We do not know what drove the government's hardline lockdown approach with all but essential workplaces being closed, but these results may have had an effect on the decision. Going forward it may have an influence decisions on extending the duration and intensity of subsequent actions.

The Prime Minister has said that what prompted her to go quickly to the level four intervention was a report showing some apparently extreme outcomes. The report came from an investment firm with limited background or expertise in the field. The content of that report has not been disclosed. It should be.

Like any modeling the OCRG results depend on the most critical assumptions used by the modelers. To a degree modeling results can be what the modellers want them to be, so it is always critical that modellers clearly report their main assumptions and their impact on the results, upfront. They should not be hidden in the technical detail.

We found that OSRG's model run results grossly overstated the number of deaths, because they made an assumption about the critical tool in the Ministry's arsenal. It was assumed that there would be no tracing and isolation of cases, or alternatively that the Ministry was so ineffective that their efforts could be disregarded. This led to an explosion in the number of cases and deaths. The reporting of the range of deaths was also inflated by the simple expedient of excluding the model runs that produced low numbers. One of their six scenarios showed just seven deaths over a year.

But that is not what the public saw. The 31 March Stuff report on the modeling read as follows:

Up to 14,000 New Zealanders could die if coronavirus spread is uncontrolled, according to new modelling by the University of Otago, Wellington.

And it could be worse, the OCRG interviewee explained:

"*Our modelling was based the hospitals and ICUs being able to function ... The death rate would certainly increase dramatically if hospitals and ICUs were overwhelmed,"*

When we ran the Covidsim model we found credible paths that could reduce the pace of infections to sustainable levels. Deaths in the range of 50 - 500 over a year are more realistic numbers. 500 deaths is around average for the seasonal flu. We found that the higher OCRG numbers were mostly generated by their assumption that tracing and testing would be abandoned.

This OCRG assumption is almost incomprehensible, unless there was a deliberate attempt to blow up the numbers. Whether the Ministry was 'in on it', or simply didn't understand what was being reported to them, we do not know. We have attempted to discuss the issue with the OCRG but have had no response.

The main purpose of this report is to look behind the ‘headlines’ to assess the robustness of the ORSG modeling, and to present our own modeling results, using the same model. We also present the arguments for the lockdown being less restrictive and present a rough cost benefit analysis of the decision to lockdown the building and construction industry.

Recently, another report has made the headlines. Te Pūnaha Matatini, a cluster of researchers from Landcare Research and the Universities of Auckland and Canterbury, has released a report that purports to show that the level four intervention was justified and suggests that it be extended. Otherwise, there would be strong growth in cases. However, the report only looked at a comparison of level’s two and four interventions. It did not consider the real question of whether some variant of a level three would be sufficient to contain the epidemic. The model can be assessed against recent actual case number outcomes. It appears to have failed.

We briefly review this modelling in this report.

Our report is structured as follows:

- Part two presents our key conclusions;
- Part three examines the Covidsim model and assesses the OSRG’s modeling outputs;
- Part four presents the results of our modeling using the Covidsim calculator;
- Part five discusses recent New Zealand and other evidence and the implications for the model’s calibration;
- Part six makes the case for relaxing certain work lockdowns using the example of the building industry. A cost benefit assessment is made;
- Part seven discusses the features of a forthcoming Tailrisk Economics model that addresses some of the weaknesses in the Covidsim model and extends the analysis to a consideration of different intervention options.
- Part eight briefly discusses the Te Pūnaha Matatini report.

Part two: Key conclusions

Ministry's modelling exaggerated the coronavirus impact by ommitting key tool
The Minstry engaged the Otago University's Covid-19 Research Group to model the impact if efforts to eliminate the Covid-19 virus failed. The research group used a simple online calculator, but failed to use a key feature of the model that allows for contact tracing, testing and isolation. This implied that the Ministry was abandoning its trace, test and isolate strategy. This has the effect of blowing out the numbers, and was largely responsible for the estimates of between 8600 and 14400 deaths if the initial eradication policies failed. These numbers may have played a role in the Government adopting a tough lockdown policy. The same model, configured with effective tracing and isolation, and some other plausible assumptions, generated about 160 deaths.

Focus of deaths needs to be supplemented with an adjustment for life years lost
Not all deaths have the same social cost. The death of a 90 year old can be sad, but the death of a child or young adult is almost always a tragedy. Burden of disease estimates often adjust for the number of life years lost and this adjustment should be made in assessments of the benefits of intervention options.

All of New Zealands nine deaths have been over 70 years old and had underlying health issues. This in line with international experience, which suggests that 85-90 percent of deaths have been in the 70+ age group.

The true burden of the epidemic can be calculated by applying an factor of around 0.15 to the number of account for life years lost. 500 deaths becomes, 75 on an adjusted basis, and can be compared with the 350 lives lost on the roads each year.

The right balance of testing and isolation and social distancing essential
Effective tracing, testing and isolation, and broader social distancing policies are both essential, but at the margin can be substitutes. The more effective the case isolation process, the less the need for costly workplace closures.

New cost benefit model addresses issues with Ministry's modelling
Tailrisk Economics has built a new model that addresses issues with the Covidsim model used by the Otago University Covid-19 Research Group. It can flexibly assess the impact of changing interentions over time; distinguishes between the high risk group and others, and makes cost benefit assessments of policy interventions.

The decision to lockdown all non-essential work places not justified

The decision to lockdown all 'non-essential' workplaces was an overaction that was not supported by evidence. Many workplaces pose a low risk of spreading the virus, but their closure can be economically costly. Our assessment of the closedown of the building and construction industry for one month is that the costs were around the $3 billion, and the benefits in terms reductions in lives losts, illnesses and hospitalisations over the next year,were $7.6 million. One life was saved.

Full disclosure

The Ministry should release all of the modelling documents that may have contributed to policy formulation.

Part three: What drives the covidsim model and the OCRG results

How the Covidsim model works

The Covidsim model is an interactive online calculator that can be run by any reasonably informed modeller, and can produce results within minutes. However, gaining a deeper understanding of the models properties and limitations takes considerably longer.

The heart of what drives the model results is the assumption on the what is called the basic reproduction rate 'R_0'and how this is impacted by policy interventions. R_0 describes how many people an infected person will infects over the average infectivity period, which in this case is about 10 to 11 days. This number is difficult to assess from the data in the early stages of an epidemic, but is assumed to be between two and three by many modellers. A rate of, say 2.5 means that over the 10 or 11 day infectivity period, the average case will pass it on to 2.5 people, assuming no change in voluntary or enforced behaviour. It is the do nothing assumption. The simple maths shows that from a small starting base of infections, a large part of the population will become infected over three or four months or so. Ten cases will increase to nearly 40,000 in three months. This reflects the power of exponential growth.

The basic reproduction rate will differ from society to society. Densely populated urban areas (New York for example) will have a higher R_0 than a society of hermits. New Zealand probably sits on the lower end of the scale, with its less densely populated cities, less reliance on crowded public transport, and, possibly, less physical contact in social greetings. A mumbled 'gidday Trev' is less risky than six kisses. Kiwi reserve has its value. There is case evidence for this. About 45 percent of reported cluster case numbers in New Zealand appear to relate to events where alcohol played a part in reducing social distance and inhibitions.

The OCRG discussed the R_0 calibration in its first report to the Ministry dated 24 February, which looked at the impact of a completely unconstrained covid-19 epidemic in New Zealand. They settled on a R_0 figure of 2.3, but made the case for adopting a lower number, given the some of the above factors.

The purpose of policy interventions is to reduce R_0 below that would spontaneously occur as the population reacts to the risk posed by the virus. If R_0 can be reduced to below one, the epidemic will eventually die out- absent any external injection of new cases. Any R_0 above 1 will eventually infect most of the population, but the time it takes will be longer, depending on how far above one R_0 is. However, a figure above, but close to one, means that peak pressures on the health system can be reduced, and a vaccine might come to the rescue before most of the population is infected.

Model interventions

There are four interventions in the model.

1. **Border restrictions**

   The model has a setting for the number of infected people arriving each day from overseas. Precautionary measures are assumed in the OCRG modelling to be almost completely effective, with one new case each day. We have used the same number in our base case modelling.

2. **Case isolation**

   Case tracing and isolation is captured by the probability that a case will be identified and isolated. In the OCRG modelling this option is not used. There is no case isolation. In our view this is a serious error because case isolation is at the heart of New Zealand's epidemic control strategy. The OCRG's failure to model case isolation appears to account for much of the modelled deaths. We have assumed an effectiveness rate of 60 percent. This is probably higher than what the Ministry is currently achieving, but we have allowed for a moderate improvement over time.

3. Isolation of sick cases

If a case is hospitalised then it assumed that there is zero transmission back into the community. This assumption, which is optimistic, is hardwired in the model. However, once the number of hospital isolation beds are full, the sick have to recover at home, where there is a probability of the infection being passed on.

The policy instrument here is the number of hospital isolation beds. The OCRG assumes only 1500, or 10 percent of hospital beds. We have assumed that this capacity could be rapidly expanded as necessary, and have assumed a capacity of 15000. Hotels and hostels, and so on, can be readily converted to quarantine stations. Some countries that have run successful case management strategies have required detention though the course of the illness, and some have not. It turns out that in our more realistic model runs this expanded capacity of 15,000 is not needed. In New Zealand only about 2-3 percent of cases have been hospitalised, so we are largely taking the nonquarantine approach.

The second variable here is the effectiveness of home isolation. The OCRG assumes 50 percent effectiveness. There is no discussion on why this figure was selected. In combination with the low 'quarantine' capacity, this low effectiveness number can set up an explosive growth in case numbers. We think New Zealanders will largely play by the rules, and have assumed that home isolation is 80 percent effective. The Ministry's strategy of home isolation seems to reflect this higher degree of confidence.

Unusually, the Covidsim model does not allow for a relationship between hospital capacity (especially ICU and ventitator capacity) being exceeded and death rates. This is a central feature of some models. Instead the focus, in terms of death outcomes, is on quarantine capacity rather than medical capacity.

4. General contact reduction

There is a single model input reflecting the intensiveness of measures designed to reduce the level of interactions in the community, and so the effective rate of transmission. The OCRG assumes reductions of 50 and 25 percent in two sets of model scenarios. These numbers are low (the full lockdown reduction in interactions is probably more like 70 to 80 percent), because the OCRG were focussing on interventions that are sustainable for a long period of time. There is no explanation, however, of why those particular numbers were selected.

The 25 percent reduction is probably consistent with almost no government intervention, beyond advisory statements and promotions, and perhaps a ban on larger public gatherings. People, particularly the vulnerable (who matter most for the death outcomes), have reduced their interactions spontaneously, motivated by self-protection and public spiritedness. The 50 percent assumes some level of imposed constraints.

We have selected a 35 percent reduction implying a moderate average level of official intervention over the full one year modelling horizon.

5. **Length of the intervention**

The length of the interventions can be varied, but once the intervention is taken off, social interactions are assumed to revert to their normal pre-crisis levels. The OCRG assumed intervention periods of six months and nine months in their modelling. The problem with the six month period, in particular, is that once the intervention is taken off the model will revert to an explosive growth cycle again, as there is nothing to contain it. This significantly impacted on deaths over the year.

We have assumed that contact reductions are sustained over the whole year. It is now generally accepted that there will be no return to 'normal' and that some behavioural changes and imposed restrictions will be with us until, hopefully, a vaccine comes to the rescue. Our 35 percent general contact reduction assumption reflects that reality.

The death rate

The model generates a number of those infected and the number of deaths will depend on the death rate/infection assumption. The OCRG assumed a death rate of about 0.7 percent, after a careful consideration of the relevant evidence. We have used that assumption.

Summary

General and case management strategies are partial substitutes, particularly in the earlier stages of an epidemic. The more intensive and effective the case management, the less intrusive and costly the general contact restrictions need to be.

At heart the Covidsim model is simple. A R_0 is selected and that is reduced by the interventions to an effective rate. If we start with a R_0 of, say, 2.5, and the general and hospital interventions between them reduce it by 40 percent, and case management by a further 40 percent, then the effective R_0 falls to 0.9. The epidemic will tend to peter out, with a low level of infections being sustained by the continued small number of new cases from abroad. If, on the other hand, those interventions are less effective with reductions of 35 percent each, the effective R_0 will be 1.06 and cases will grow over time. As Mr. Micawber would have put it. 'R_0 equals 0.9, result happiness. R0 equals 1.1, result misery.'

The OCRG modelling

The OCRG ran six scenarios with the Covidsim model. Three R0 assumptions (1.5, 2.5 and 3.5) were combined with two intervention scenarios that assumed either a 50 percent contact reduction in R_0 over six months; or a 25 percent reduction over nine months. As previously noted there was no reduction in the post intervention effective R_0 reduction due to case management.

The results are shown in table 1.

Table 1: OCRG scenario results

Model assumptions	Deaths over one year
R0 1.5 25% control over 6 months	2520
R0 1.5 50% control over 9 months	7
R0 2.5 25% control over 6 months	12,700
R0 2.5 50% control over 9 months	8560
R0 3.5 25% control over 6 months	14400
R0 3.5 50% control over 9 months	11800

As noted above, the R0 1.5 results were not reported in the media. The range was reported as between 8560 and 14400. If the OCRG had little confidence in the 1.5 estimate, then they should have replaced it with a more plausible lower estimate, such as a R_0 of 2, and then reported that number. Similarly the upper estimate could have been set at a high, but still reasonably possible 3. Instead the reader is given a range of between 8560 and 14400 deaths, giving the misleading impression that there is a good deal of certainty around the estimates of high death numbers, because the upper and lower estimates are relatively close together.

More importantly, the 'headline' results are absent the impact of any case management. The results should never have been released without explaining that there was no contract tracing and isolation process in effect. The numbers with contract tracing should also have been prominently reported.

Issues with the Covidsim model

In the main we found the Covidsim model to be intuitive and (mostly) easy to use. However, there are some weaknesses in the model, and some enhancements would contribute to its usefulness.

1. Time varying intervention inputs

At present the interventions can only apply for a set period, and reverts to zero if removed. There is no capacity to set them at a higher level or lower, non-zero, level over the course of the one year model run. Time varying intervention coefficients would allow the user to easily alter the intensity of the interventions over time, better reflecting how policy is likely to be run.

2. **An array of general contact reductions by the intensity of the intervention**

At present all of the impacts are subsumed in the single number. This makes it more difficult for the outsider to assess the impact of intervention changes. An array of inputs could be presented showing the relationship between contact reductions and the policy actions. The array could look something like those set out in table 2. The table is purely illustrative and does not reflect a final view of the content of the table, or of the calibration of the contact reductions.

An alternative would be to link the contact reduction with the Government's intervention levels. The problem here is that these levels are insufficiently sensitive to the intensity of the intervention. Level three, in particular, can come in a variety of levels of intensity, possibly ranging from contact reductions of 30 to 60 percent. And all lockdowns are not equal particulary with respect to their costs.

Table 2:Inteventions and interaction reductions

Intervention	Reduction in interactions %
1. General social distancing behavioural changes	25
2. Above plus work at home where possible bans on high risk activities (weddings and social events) above a certain size.	35
3.Above plus ban on wider range of events	45
4. Above plus restaurant shut down	60
5. Above plus shut down of other higher risk businesses	70
6. Above plus shutdown of all 'nonessential 'businesses	75
7. 'Wuhan' style	80

3. **There is no consideration of the costs and benefits of the intervention strategies.**
   This was not the Covidsim modellers', or the OCHG's mandate. But it is the Health Ministry's and Treasury's. Hopefully this has already been done, or is underway. If not it should be done quickly. Part six explores some of the issues involved in getting the building and construction industry back to work, using a cost benefit approach.

4. **It does not distinguish between the high risk group and the general population**
   Separately identifying the high risk group (over 70's and others with existing conditions) would allow an analysis of the impact of applying different general contact reductions to both groups. The non at-risk group could have a contact reduction ratio of 35 percent, but it could be increased to, say, 70 percent for the at risk group. It would be useful to allow for an interaction between the two variables. An increase in the non at-risk group contacts would generate a higher at-risk effective R0 for the at risk population, capturing the leakage from one group to the other.

4. **The headline death number needs to be supplemented by a life years saved metric**
   The headlines death number implicitly assumes that all lives lost are valued the same. But that is not how society operates. The death of an already ill 85 year old is not regarded as equivalent to the death of 17 year olds with their lives in front of them. In an average year 500 people die in our seasonal flu epidemics, but there are no calls for the economy to be shut down to extend the lives of the elderly who account for the bulk of the deaths. Every death does not receive headline reporting in the media.

   Health researchers and economists adjust deaths by the number of life years lost, to provide a truer picture of the burden of the illness, and the value of the measures taken to reducue that burden. If the average years lost is 15 percent of those lost in an average road accident, then the age adjusted number of lives lost would be one fifth of the headline number. So 500 lives lost becomes a life year adjusted 75 lives lost, compared to the 350 lost in a year on New Zealand's roads.

   The distinction between deaths and life years lost is particularly important because 85-90 percent of deaths world wide have been in the 70 plus age group. Most would have had underlying medical conditions, and would have had a limited life expectancy. Most of the New Zealand deaths, so far, have been in nursing homes where the average stay is two years.

The life years lost adjustment would provide a different view on the social costs if the containment measures do not work. In its March 24 report the OCRG presented the following perspective on its worse case outcome of 27,600 deaths.

This death toll would far exceed the death toll for NZ from World War One (18,000 deaths) and from the 1918 influenza pandemic (9000 deaths).

If we adjust for the difference in population size (4.5 times) and life years by applying an adjustment factor of 0.15 (the 1918 epidemic dispropotionately affected younger adults) then the 27600 shrinks to 915, significantly below the 1918 pandemic disease burden.

Part four: Our modeling results

In this part we set out the result of model runs using more realistic assumptions than in the OCRG modeling. In particular we do not assume that there is no contact tracing and isolation. Table 3 summarises the key assumptions that were discussed in part three. Our benchmark model run shows 105 deaths after six months, and 157 after one year. This is broadly consistent with the experiences countries such as South Korea, Hong Kong, Taiwan, which have achieved a good measure of control over their epidemics without the need for harsh lockdowns.

In table 4 we present the results of our sensitivity testing. We start with our base case, and then alter the key variables, one by one, to understand the sensitivity of the results to changes in each variable. For example, if the initial infections are reduced from 500 to 100 then the number of deaths over the year fall from 157 to 53. Sensitivity testing is essential to understand which variables matter the most, and also gives insights into the properties of the model. To assess the impact of a particular assumption, simply compare the six month and one year results in table 4 with the base case model results.

Table 3: Summary of input variable assumptions

Variable	**Assumption base case 1**	**Base case 2. As for 1 except where indicated**
Imported cases	1 per day	
Initial infections	500	
Contagiousness R0	2.25	3
General contact reduction	35%	45%
Length of general reduction	365 days	
Case isolation - identification	60%	
No of isolation beds	30 per 10,000 population	Nil
At home isolation effectiveness	80%	
Length of isolation period	365 days	
Seasonal fluctuation in infections	Plus 25%	
Infections lead to sickness	67%	
Sick patients hospitalised	5%	4%
Sick patients die	0.7%	

Table 4: Sensitivity analysis

	Cumulative deaths at 182 days	**Cumulative deaths at 365 days**
Base case outcome	105	157
Initial infections (base case 500)		
100	33	53
1000	193	281
Imported infections (1)		
0.5	98	144
2	119	182
10	235	382
Contagiousness (2.25)		
2.5	400	865
2	34	41
3	6108	11309
Case isolation		
Cases identified % (60)		

0%	15523	16197
50%	346	725
70%	38	46
At home isolation of patients (80) effectiveness		
50%	105	157
At home isolation effectiveness plus no hospital isolation		
50%	3414	7317
At home no hospital isolation base case 2 80% home effectiveness	443	990
General effectiveness (0.35)		
0.3	262	505
0.50	15	17
Length of general measures (365)		
182 days	105	973
Death rate (0.7)		
1.4	210	313
Base case 2	988	2256

Sensitivity testing commentary

All of the assumptions have an impact on the death outcomes, but some matter more than others.

Case isolation

Case isolation is critical. With 60 percent effectiveness, the number of deaths is 105 at six months and 157 over the full year. With no case isolation tracing, testing and isolation, deaths blow out to 15,500 after six months. It is not possible to contain the epidemic without the most extreme and sustained general population contact suppression measures.

Contagiousness

The results are very sensitive to the contagiousness assumption. Our best estimate contagiousness assumption (R_0) of 2.25 generates 157 deaths. An R_0 of 2.5 generates 865 deaths.

Length of intervention

Removing all interventions after six months has a material impact. Cumulative deaths by year-end increase from 157 to 973.

Starting point

The results are not as strongly dependent on the starting number of undiscovered cases. The proportionate increase in deaths is less than the proportionate increase in the starting point number.

At home isolation

The results appear to show that deaths are insensitive to the rate of compliance of sick patients with isolation requirements. This is because the model is calibrated so that all of the isolation is done in hospitals. If all illnesses are assumed to be treated at home (close to New Zealand practice), then there is a strong sensitivity to this isolation assumption.

Death rate

Deaths have a linear relationship with the death rate assumption.

Overseas travel

The rate of new infections from overseas travel is not so important, given the implied assumptions of limited overseas travel and strict border controls. Even if this situation changes, and 10 cases a day slip though border controls, compared to one for the base case, deaths only increase from 157 to 382 over the year. Here technological developments could make a difference to the opening up of international travel. If every passenger was given a 15 minute Covid-19 test before boarding the plane then the risk from international travel could be reduced to a minimal level. Alternatively travellers could be tested on entry, with positive cases quarantined.

Part five: How are we doing?

It is still early days in New Zealand's epidemic, but the numbers do provide some useful information for the calibration of a model. Some of the data has less direct relevance, but is still interesting.

Figure 1 below shows the daily count of new cases taken from the Ministry's website. There has been a leveling out, and then a fall in the number of cases, which would have been more than twice as high if the previously trajectory had been maintained. If is likely that most of the 'flattening of the curve' in the 10 days or so subsequent to the lockdown, would have been due to pre-lockdown measures and voluntary behavioral changes, because there is a material lag from changes in behavior to case identification.

Figure 1: New Zealand daily cases as at 14 April

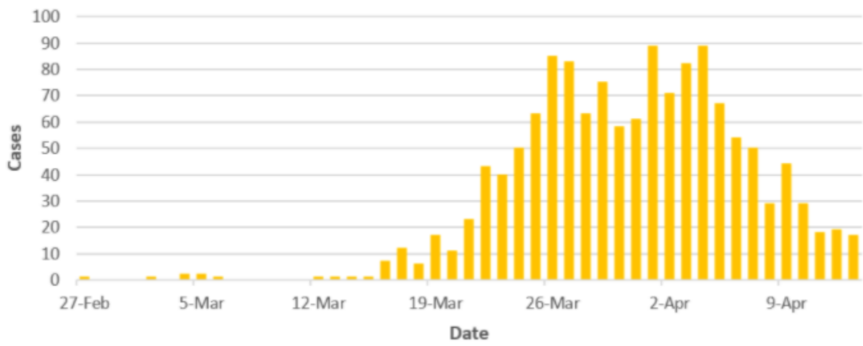

New confirmed and probable cases over time, as at 9.00 am, 14 April 2020

Source: MOH 14 April

Deaths

There have been 9 deaths up to April 15. All were over 70, and three were in their 90s. This is in keeping with other countries where around 85 to 90 percent of the dead were over 70 or older. Deaths of under-40 year olds without underlying medical conditions are very rare.

Other important take-outs include:

- Hospitalisation rates have been low. As at 14 April there were 15 people in hospitals and three in critical care. The cumulative number of hospitalisations is not now disclosed (it should be) but from previous information provided it appears that the hospitalisation rate is about 2.5 percent.
- The age distribution of cases appears to be favourable. Seven percent are in the 70+ age group. This is lower than the share in the population (12 percent), and is

materially lower than countries that experience the highest numbers of deaths. In Italy 36 percent of cases (Statistica 6 April) were over 70. Nearly 50 percent of New Zealand cases were under 40 and were not exposed to a significant mortality risk.

Figure 2: Cases by age group

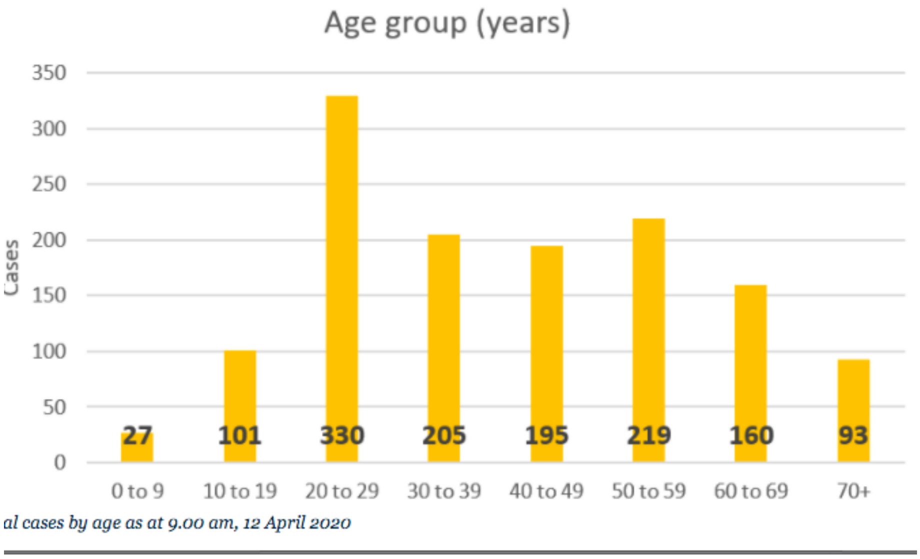

Source: MOH

- Maori, who are also identified as a vulnerable group, account for around 8.5 percent (Ministry 11 April) of cases compared to a population share of 15 percent.

Two weddings and no funeral

The 'cluster' data provides less information than it did previously. Clusters are now only reported when they exceed 10 cases. A lower limit would provide more information. It is clear that social occasions have been important sources of contagion, accounting for 45 percent of cluster numbers. Two weddings have resulted in 98 cases (11 April). There is no reported funeral cluster.

There was one reported workplace cluster of 28, but we do not know whether it is a real workplace or not. At the Bluff wedding none of the staff were infected. 'Workplace' has now been replaced by the descriptor 'community' so it is not clear what is happening. The limited information we have suggests that workplaces have not been a substantial driver of contagion, possibly accounting for about 5 percent of cases.

Undetected community spread

It is difficult to get a fix on the amount of undetected community spread. The Ministry's website says 2 percent of cases were community spread – with a further 11 percent as yet unidentified. This low share suggests that casual contact, for example, in a supermarket, is not a significant source of contagion, so much of the social distancing measurers may not be reducing risk by very much, at least at this stage of the epidemic. It may have value, however, in getting people habituated to different behaviours, which may be necessary for some time.

The identified share of community spread has not been growing, which is reassuring. However, the Ministry has not being doing much community testing, so the 2 percent may not be very meaningful. One other piece of assurance is the proportion of cases to the number of tests. If this is low, even when the Ministry has been doing contact tracing, where the probability on a positive is much higher than in the general population, then this may be a good sign.

Table 5 shows this ratio for a range of countries. New Zealand's average of 3 percent is a good rate, though, again, Australia appears to being better at 1.9 percent, reflecting their more intensive testing. Recently the New Zealand ratio has been trending down. Over the last three reporting periods it was about one percent. In Iceland, the world's most intensive tester, the ratio is still over two percent.

Table 5: Proportion of positive tests

	Cases '000	Tests '000	Prop.
US	312	1633	0.191
Spain	126.2	355	0.355
Italy	124.6	657	0.190
Germany	96.1	918.5	0.105
France	89.7	224.2	0.400
China	81.7	NA	NA
Iran	55.7	186	0.299
UK	41.9	183.2	0.229
Turkey	23	161.4	0.143
Switzerland	20.5	163.4	0.125
Belgium	18.4	70	0.263
Netherlands	16.6	75.4	0.220
Canada	13.9	318	0.044
Austria	11.8	100.1	0.118
Portugal	10.5	81.1	0.129

S. Korea	10.2	455	0.022
Israel	7.9	90.4	0.087
Sweden	6.4	39.9	0.160
Australia	5.6	291	0.019
New Zealand	**1.0**	**33.1**	**0.031**

Source: Worldomater

Table 6 presents the Health Board case data in terms of cases per million of population, which is a common metric of comparison.

Table 6: Population adjusted case rates

Health Board	Population (000)	Cases	Cases/million
Auckland	564	180	319
Counties Manukau	579	103	177
Northland	182	26	138
Waitemata	647	200	309
Bay of Plenty	243	41	202
Tarawhiti	49	3	60
Lakes	111	16	144
Taranaki	121	14	116
Waikato	428	177	414
Capital and Coast	322	88	273
Hawkes Bay	167	41	245
Hutt Valley	151	20	132
Mid Central	180	28	155
Wairarapa	45	8	178
Whanganui	65	7	108
Canterbury	582	139	148
Nelson Marlborough	152	48	254
South Canterbury	61	12	196
Southern	334	210	628
West Coast	32	6	188

Sources: Statistics NZ, M0H 15 April

Some relevant overseas experiences

Figure 3 shows the Australian data. Australia (our most relevant comparator) appears to be doing better than New Zealand in terms of the improvement in the trajectory of new cases. It is may be a little early to say whether this is statistical noise, or says something more fundamental about the respective approaches.

However, the data appears to suggest that New Zealand's more aggressive approach on employment (Australia allows travel to work when it is not possible to work at home) has not secured any material benefits. While Australia has taken a less aggressive approach on workplace shutdowns, it has had more testing. Their cumulative testing rate per million of population was 11,700 (Worldometer 6 April), compared to New Zealand's 7600. New Zealand has now largely caught up but the Australian's earlier more intensive approach, may have contributed to their success in reducing case numbers.

Figure 3: New cases Australia to April 13

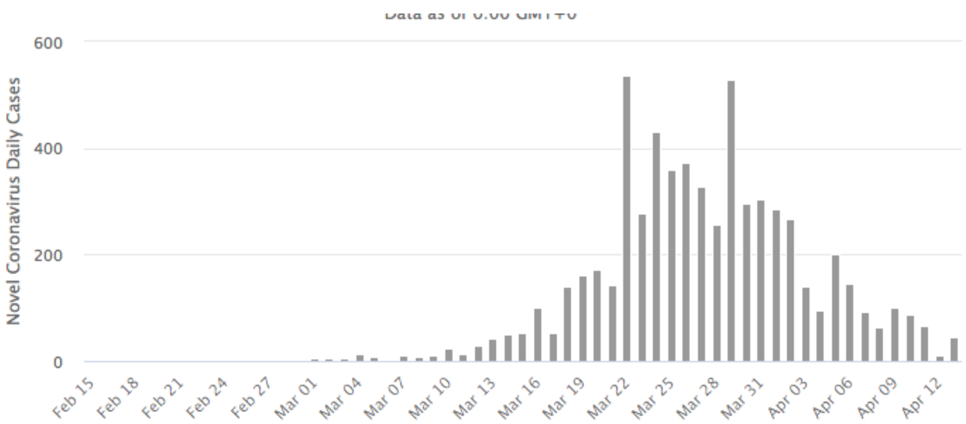

Source: Worldometer

Another relevant comparison is with Sweden, which is a stand out in its reluctance to take early aggressive measures. Sweden appears to be at level two, and has partially embraced the herd immunity approach. They appear to have looked at their numbers and come to the view that giving some of the elderly six months to a couple of years more life (possibly in a miserable lockdown for the best part of it) is not worth the devastating economic and social costs of a full blown eradication campaign. Better to recommend that the vulnerable take enhanced measures to protect themselves.

Figure 4: Sweden daily new cases to 13 April

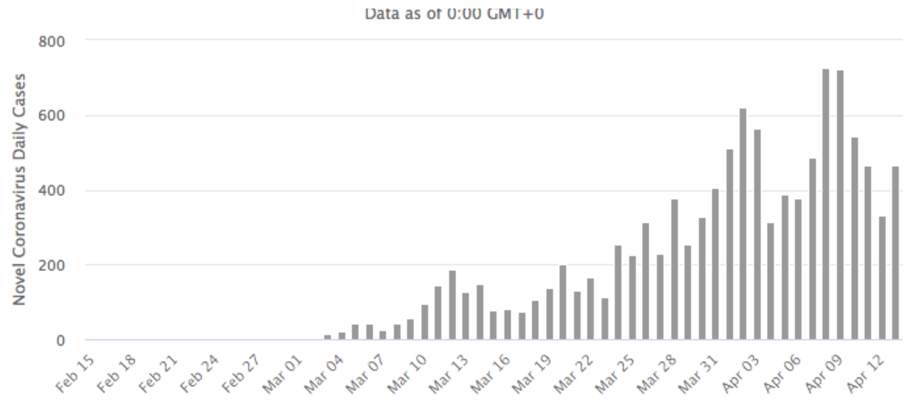

Source: Worldometer 7 April

Iceland is a standout in its approach to testing and strict isolation of positive cases. Over ten percent of the population has been tested compared to less that one percent for New Zealand. They have relatively few mandatory social distancing requirements.

Figure 5: Iceland new cases to 13 April

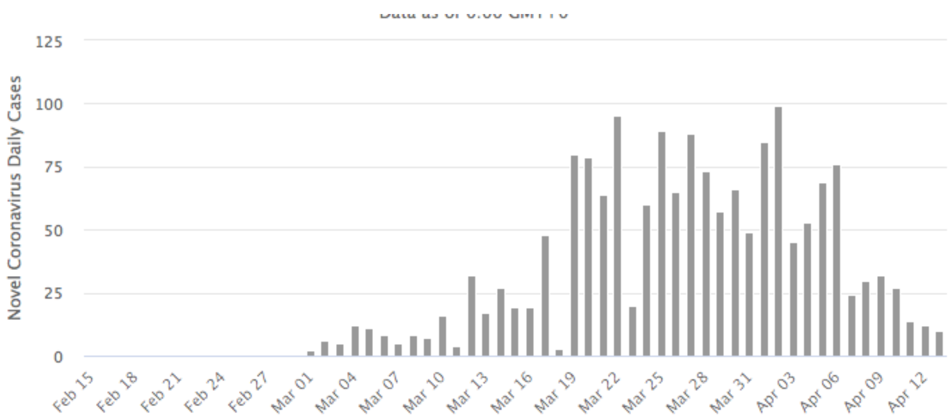

Part six: Opening up the construction and building industry

We think that there is a strong case for putting the building and construction industry back to work now, and that it should never have been closed down. Similar cases could be made for getting back to work in other areas where it is not possible to work from home, but we have focused on just this sector because of its size, and to illustrate how the costs and benefits could be calculated. Our calculations focus on the costs and benefits of excluding the industry from the initial shutdown

The arguments

The logic for opening up the industry is that:

- The workforce is most younger and well, and at lower risk of direct infection. If they are infected severe outcomes are rare.
- Much of the work is outside, in relatively small groups, so there is a lower risk of worker-to -worker infection.
- Even if there is cross infection at work the consequences are limited in terms of further infections, with the household bubbles and other social distancing measures.
- Other jurisdictions, for example Alaska and Belgium with otherwise strict lockdowns have taken this approach.
- There appears to be no evidence or analysis anywhere supporting a lockdown for the building and construction industry as a necessary part of a containment strategy, or even an eradication strategy.
- There is no evidence of the food industry, which employs a lot of workers, often with large numbers in indoor environments, experiencing any cluster event in New Zealand.
- Spain has just announced that its construction industry is going back to work, even though their daily case number, on a population adjusted basis, is about 300 compared to New Zealand's 15.

The costs and benefits

The costs

Assuming 250,000 workers are affected and that the average cost is, say $3000 a week per worker, a months shutdown will cost about $3 billion.

The benefits
To calculate the benefits we used some preliminary outputs from our forthcoming CORONA-1 cost benefit model (see next part for more detail). Note the modelling is preliminary and the results are presented here to give a sense of the possible magnitudes.

We went through the following steps.

1. We assumed that the full lockdown reduces contacts by 75 percent. We ran this level for 30 days and then reverted to a contact reduction of 35 percent for the rest of the year. The following outputs were generated.
    - The number of deaths for the vulnerable and general populations;
    - The number of illnesses and hospitalisations;
    - The cost of the deaths, hospitalisations and illnesses.

2. The impact of excluding the building and construction industry was assumed to reduce the contact reduction by 5 percent to 71.25 percent. The model was run with this setting for 30 days and then reverted to a contact reduction of 35 percent. The same outputs for the full lockdown were generated in model run 2.

3. The results from model run 2 were deducted from run 1 to produce the marginal benefits of the building and construction industry lockdown.

It was assumed that:

- The value of a statistical life is $4.5 million;
- The life years conversion factor is 0.10 for over 70s and 0.55 for under 70s;
- The cost of an illness is $4000;
- The cost of a hospitalisation is $30,000.

Results

The results are shown in table 7.

Table 7: Benefits of building and construction lockdown

Outcome	Number or value over one year
Number of deaths avoided–over 70	1
Number of deaths avoided – under 70	0
Total	1
Illnesses avoided	1400
Hospitalisations avoided	14
Benefit of Deaths avoided $'m	1.6[1]
Benefit of hospitalisations avoided $'000	0.4
Benefit of illnesses avoided $'000	5.6
Total benefits	**$7.6 m**
Total cost	$3,000 million
Benefit cost ratio	.003

Of course, the benefit cost ratio of .003 is from just model run. Different, and plausible, assumptions can readily generate benefits that are a order of magnitude, say, ten or twenty times, higher than the $7.6 million. But it is very difficult to see how they could be over 300 times higher.

Part seven: A new model- CORONA-1

To address some of the issues identified with the covidsim model we have developed a new model that builds on the same conceptual framework. The title CORONA is an acronym for Calculate Objective Results Of Nimble Actions. Corona is also a beer. Our risk models are typically named after beers.

1 This is the value of statistical lives saved, including partial lives, not full lives as reported above.

The first model will be a 'stripped-down' version that focusses on the impacts of policy interventions and their social and and economic costs and benefits. Later models could have enhancements that focus, for example, on hospital service constraints and their impact on deaths and other outcomes.

The key enhancements compared to the Covidsim model:

- *Time varying policies*
  This allows the user to change the intensity of general policy actions, and to adjust case management effectiveness over time. We see this as a critical enhancement. Changing policy settings over time as information on policy effectiveness emerges, is an important part of any management strategy.

- *Differentiation between vulnerable (over 70) and the general population.*
  This allows the user to apply different general population effectiveness parameters to the vulnerable population, change other parameters, and to report separate outcomes.

- *Social costs of the epidemic.*
  This translates death and illness outcomes to social costs for both the vulnerable and general populations. This recognises the different social costs of deaths in different age groups by assessing the impact in terms of life years lost.

- *Economic and social costs of the interventions*
  This quantifies the economic and social costs of a range of intervention initiatives and allows a calculation of the costs and benefits of various intervention strategies.

Part eight: The Te Pūnaha Matatini report

The Te Pūnaha Matatini news event

Recently the key results of a paper from Te Pūnaha Matatini was reported in the media, The headline in the Herald on 9 April was:

Coved 19 coronavirus: New data reveals bullet NZ dodged by locking down when we did

And the text of the report reads:

New modelling has revealed the bullet that New Zealand dodged by going into lockdown, with data suggesting officials would've otherwise been reporting several hundred new cases today.

Instead, the Government this afternoon reported just 29 cases – making today the fourth in a row that numbers had dropped from the day before.

The alternative scenario could have seen around 200 new cases announced today – and perhaps 350 daily cases reported around this time next week.

The story had a supporting comments from the Te Pūnaha Matatini director Professor Shaun Hendy:

'*But in our graph you can see that if we hadn't gone into lockdown, then new case numbers would have been steadily rising since the end of March.*

"The fact that the actual case numbers have not risen shows that lockdown is having an effect. Next week we should have a good sense of how much of an effect and what this means for our lockdown."

Except the Director's story was not supported by the text in the report, which said :

These measures were enacted on 25 March, and at the time of writing, there has been an insufficient period of time to fully observe their effectiveness from clinical case data alone.

The story was also supported by the following figure provided by Te Pūnaha Matatini. It shows the difference between the modelled number of cases at level 1 and those for alert level 4. But,of course, reverting to level 1 was never an option when the lockdown was announced. But there was a choice of various less intrusive levels of intervention at alert levels three and four. The impacts of most of the possible interventions were not analysed and/or reported, and the public was left with the impression that the only choice was between hundreds of new cases a day, or a strong level four response that involved shutting down most workplaces.

Figure 6: TPM Media presentation case number simulations

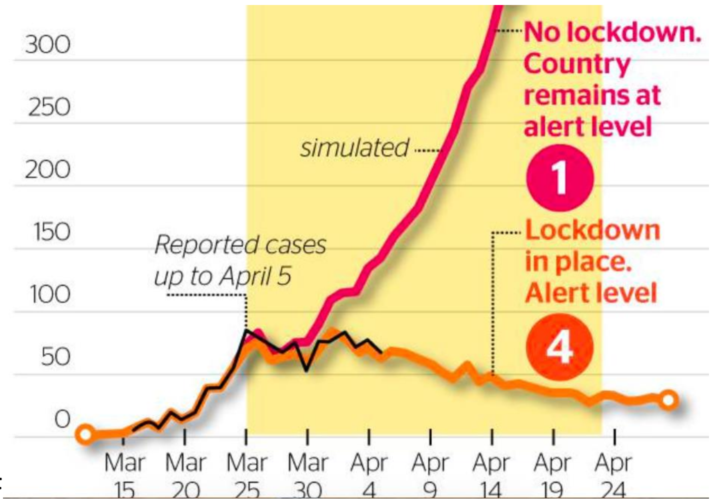

F

The TVNZ story pushed the TPH lockdown extension message. The headline read. **'*NZ can eliminate coronavirus if lockdown is extended, latest modelling shows*'**

The Te Pūnaha Matatini Model

Here we briefly describe and critique the model and the outputs which were at the centre of the Herald story.

The TPM model

It is similar in its basic structure to the Covidsim model, though it is more sophisticated in some respects. It also captures both case and general social distancing measures to generate an effective R_0, termed R-eff.

The underlying fundamental R_0 is assumed to be 2.5, which is reasonable. The overall effective R_0s, for different levels of intervention, were selected somewhat judgmentally, after a consideration of some evidence from two sources. Some was taken from an unpublished paper by a group including one of the authors of the TPM report. The evidence from the experience of 11 European countries was taken from a published paper[2]. This paper reports estimates of various interventions and their impact on R-eff, but the confidence bands were very wide, and the interventions were not statistically different from each other.

There was a table in the TPM report, reproduced below, of some illustrative country examples. The assumptions that were fed into the model are shown in table 8. The R-eff numbers drove the results. If they were above one, case numbers will grow. If below one, the epidemic will be bought under control.

The authors have assumed that a level 4 intervention generates an R-eff of 0.74 in their realistic scenarios, which will control the epidemic. The level two and three intervention R-effs, of 1.68 and 1.3, will not.

Table 8: Country examples for R-eff. by alert level

Alert	Effectiveness		
	Low	Med	High
Level 4	2.1 (e.g. GBR)	1.3-1.6 (e.g. DEU)	0.9 (e.g. NOR)
Level 3	1.8 (e.g. USA)	1.3 (e.g. NLD)	1.0-1.1 (e.g. NSW)
Level 2		1.6-1.8 (e.g. SWE)	1.1 (e.g. HKG)

Table 9: Model effective R_0s by alert level

	Optimistic			Realistic			Pessimistic		
	$C(t)$	R_c	R_{eff}	$C(t)$	R_c	R_{eff}	$C(t)$	R_c	R_{eff}
Level 1	1	2.5	2.3	1	2.5	2.3	1	2.5	2.3
Level 2	0.64	1.6	1.49	0.72	1.8	1.68	0.80	2.0	1.87
Level 3	0.44	1.1	1.03	0.52	1.3	1.21	0.60	1.5	1.40
Level 4	0.28	0.7	0.65	0.32	0.8	0.74	0.36	0.9	0.84

The results are entirely dependent on the analysis supporting the key R0-eff assumptions. We had the following issues with this analysis.

- As noted, part of the evidence supporting their judgments was not available publicly;
- The assessments appear to be biased to favour level four, over other intervention levels. For example, the realistic optimistic assumption for the level two R-eff in the table is 1.49, but their own figure for Hong Kong was 1.1. The realistic assumption for level four is 0.74 but the lowest example in the table is 0.9 for Norway.

Figure 7: Hong Kong daily cases

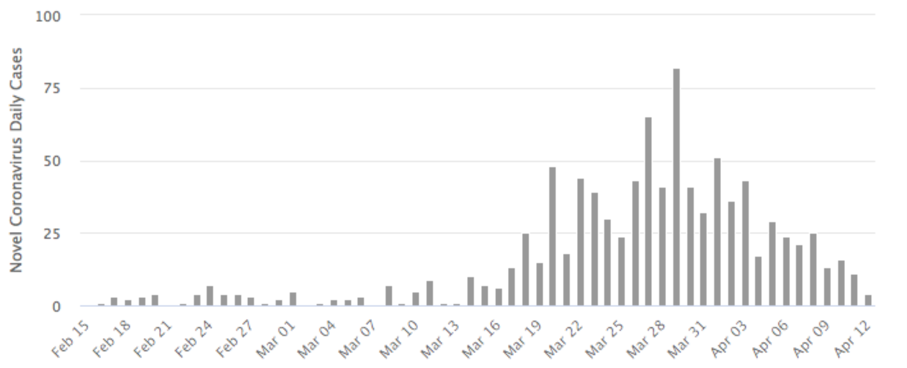

Source: Worldometer

- The estimates are out of date. While they might have been reasonable estimates a couple of weeks ago, there has been a lot of new information since then that would produce lower estimates of R-eff. For example, the Hong Kong (level two intervention) daily case count in figure 7 clearly shows that the R-eff is below one. Similarly for most of the country's discussed in part five. We could probably repeat the exercise with most of the countries in the TPM data set, if we knew what they were. Obviously, TPM could not have the most recent data in front of them when they finished their paper. But enough was available by 9 April, when the paper was released, to give them pause. Just looking at the New Zealand and Australian case numbers should have been enough.

Model outputs

There are ranges of outputs with different time horizons that look at different questions. Here we focus on just the short horizon results that grabbed the some of the headlines.

In the document there is a figure, reproduced below, that compares level two and four outcomes. Level three results must have been available but they were not shown. We have reproduced the figure below because it is easier to read than the one in the Herald report.

Figure 8: TPM Modelling paper case projections

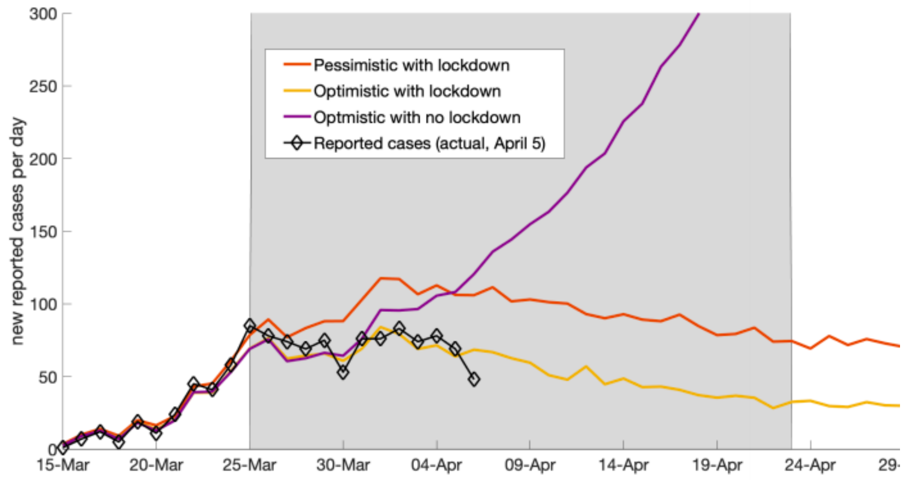

The last actual case number is shown at 5 April. At that point the level four interventions could only have had a limited impact, while increased testing could have temporarily boosted the numbers. The actual reported case number is probably a fair reflection of the level two intervention combined with voluntary adoption of social distancing. The purple line shows the estimated level two intervention case numbers at that date. It is about twice as high as the actual number, suggesting the modellers have overstated the level two R-eff. Subsequent to April 5, the actual case numbers have tracked well below even the optimistic scenario, again telling us that the model calibration is probably wrong.

Conclusion

This is a case of getting out of the model what you put in. In our view, TPM did not use the best available information, and should have either: not released their report until it was updated (and should have told a different story); or released a heavily caveated paper, without any media fanfare.