Investigation into the effectiveness of "double testing" travellers incoming to the UK for signs of COVID-19 infection

[redacted] on behalf of the PHE modelling cell

Executive summary

In this paper, Public Health England's previous work on border screening is extended to assess the effectiveness of requiring all incoming travellers to undergo two rounds of PCR type testing. One test at arrival at UK border and again some period (the time required to take, run and report a test) before ending of that person's quarantine.

The model simulates, via a Monte Carlo method, the possible epidemic and testing trajectories of individuals infected with SARS-CoV-2 at an unknown time before travel, attempting to enter the UK population and the points at which the disease may be detected. Multiple scenarios are considered defined by combination of flight time (ranging from short-haul, to long-haul) and isolation period (7, 10 and 14 days). A baseline of only performing PCR testing at arrival is also provided.

Fundamentally, this work shows that there is an exchangeability between the second test detecting infected individuals at their peak of detectability and maintenance of individuals in quarantine until such time as they are no-longer infective.

Findings

- Requiring incoming travellers to self-isolate on arrival to the UK increases the detection rate of infected travellers compared to the base case.
- The longer incoming travellers are required to self-isolate, the higher the expected detection rate of infected travellers.
- There is a natural correlation between longer flight time and detection rate, however, this correlation becomes less pronounced as the period spent self-isolating increases.

Double testing success rates (averaged across flight times considered):

Double testing scenario	Double testing success rate
Base case (testing only on arrival)	0.07
Second test administered after 5 days isolation, before being isolated for a further 2 days (isolated for a total of 7 days)	0.85
Second test administered after 8 days isolation, before being isolated for a further 2 days (isolated for a total of 10 days)	0.96
Second test administered after 10 days isolation, before being isolated for a further 4 days (isolated for a total of 14 days)	0.98

Please note that a more detailed companion piece to this work, commissioned from APHA and considering greater nuance such as individual country prevalence, is currently under rapid review and should be available next week.

Contents

Executive summary 1
Findings 1
Background 2
Results 4
Conclusion 5
Appendix 1 – Model description, assumptions and parameterisation 5
Model description 5
Modelling assumptions 7
Model Parameters 7

Background

In this work, we use stochastic modelling to assess the effectiveness of implementing a "double testing" strategy at the UK's borders. By this we mean that all incoming travellers will undergo screening on arrival into the UK, where they shall then be isolated for a certain number of days before being tested again, and if they test negative they shall be released into the UK general population.

The following work builds on previous work also done by Public Health England, which focused around the effectiveness of implementing entry screening at UK borders (currently in pre-print), where we modelled the impact of subjecting all incoming travellers to a single screening process on arrival. This work differs in that we introduce a period of isolation, after which these travellers are screened again. We shall therefore take a similar approach to the previous work and model a given number of infected individuals attempting to travel from some external country, undergo the implemented border procedures, after which they shall attempt to gain access to the UK. The outputs as presented in the following work is highly dependent on a range of assumptions, which we make clear below. As such, this work should be interpreted as indicative, rather that predictive.

In this work all individuals modelled are infected (where infection occurs prior to boarding their flight), each being allocated their own incubation period/time till detectable. These individuals will then attempt to board a flight to the UK. If successful, they will then be screened on arrival, where if testing negative they shall then undergo the "double testing" procedure before entering the general population. The results reported here then only describe the probability that infected individuals gain access into the UK given that they successfully travelled to the UK (i.e. boarded their flight), and *not* the absolute numbers of infected individuals one should expect arriving at the UK's borders. This probability is then the metric by which we assess the effectiveness of the considered double testing scenarios. We apply our model to a range of different flight times as well as mitigation scenarios implemented at the border. Flight times are assigned by randomly sampling from the ranges [3,5], [7,9] or [11,13] (emulating scenarios where travellers have taken a short, medium or long-haul flight to the UK) depending on what scenario is considered. For the double testing structure used, we start with a base case, in which we assume that aside from screening on arrival, there is no self-isolation or additional testing processes being implemented at the border; this shall then give a rough

indication as to what proportion of the infected attempting to enter the UK one could expect without the additional procedures. From this, we then consider the scenarios where all incoming travellers, in addition to being tested on arrival, are then also put into self-isolation, where they are then tested again after 5, 8, or 10 days. If these come back negative, they are then released on their 7^{th}, 10^{th} or 14^{th} day of isolation respectively (if no observable symptoms have developed in the meantime).

Results

Here, we present the output as returned from our model:

Flight time range	Isolation routine	Non-fliers	Detected, arrival	Detected, iso.	Detected, sympto.	Undetected	Double testing success rate
UNIF (3, 5)	(None, None)	60296.0 (59968.0, 60518.0)	1172.0 (1124.0, 1224.0)	0.0 (0.0, 0.0)	0.0 (0.0, 0.0)	38531.0 (38312.0, 38870.0)	0.03 (0.028, 0.031)
UNIF (9, 11)	(None, None)	60286.0 (60045.0, 60505.0)	2936.0 (2849.0, 3022.0)	0.0 (0.0, 0.0)	0.0 (0.0, 0.0)	36778.0 (36550.0, 36998.0)	0.074 (0.072, 0.076)
UNIF (15, 17)	(None, None)	60314.0 (60056.0, 60596.0)	4704.0 (4591.0, 4813.0)	0.0 (0.0, 0.0)	0.0 (0.0, 0.0)	34982.0 (34714.0, 35225.0)	0.119 (0.116, 0.121)
UNIF (3, 5)	(120, 168)	60273.0 (60029.0, 60491.0)	1172.0 (1117.0, 1224.0)	29762.0 (29581.0, 29956.0)	2390.0 (2311.0, 2458.0)	6403.0 (6273.0, 6533.0)	0.839 (0.836, 0.842)
UNIF (9, 11)	(120, 168)	60284.0 (60011.0, 60502.0)	2930.0 (2843.0, 3022.0)	28809.0 (28572.0, 29045.0)	2177.0 (2111.0, 2255.0)	5799.0 (5683.0, 5893.0)	0.854 (0.852, 0.857)
UNIF (15, 17)	(120, 168)	60296.0 (60037.0, 60541.0)	4700.0 (4588.0, 4807.0)	27767.0 (27534.0, 28009.0)	1976.0 (1896.0, 2053.0)	5261.0 (5153.0, 5394.0)	0.867 (0.865, 0.87)
UNIF (3, 5)	(192, 240)	60289.0 (60055.0, 60553.0)	1171.0 (1126.0, 1223.0)	35858.0 (35638.0, 36064.0)	731.0 (691.0, 779.0)	1951.0 (1893.0, 2022.0)	0.951 (0.949, 0.952)
UNIF (9, 11)	(192, 240)	60286.0 (60060.0, 60525.0)	2932.0 (2852.0, 3020.0)	34354.0 (34150.0, 34559.0)	659.0 (621.0, 707.0)	1769.0 (1701.0, 1845.0)	0.955 (0.954, 0.957)
UNIF (15, 17)	(192, 240)	60289.0 (60071.0, 60522.0)	4696.0 (4591.0, 4805.0)	32811.0 (32615.0, 33047.0)	595.0 (560.0, 633.0)	1608.0 (1549.0, 1669.0)	0.959 (0.958, 0.961)
UNIF (3, 5)	(240, 336)	60294.0 (60067.0, 60513.0)	1176.0 (1119.0, 1231.0)	37303.0 (37051.0, 37512.0)	478.0 (445.0, 511.0)	749.0 (700.0, 782.0)	0.981 (0.98, 0.982)
UNIF (9, 11)	(240, 336)	60272.0 (60016.0, 60544.0)	2939.0 (2864.0, 3027.0)	35682.0 (35448.0, 35892.0)	429.0 (402.0, 461.0)	678.0 (644.0, 713.0)	0.983 (0.982, 0.984)
UNIF (15, 17)	(240, 336)	60287.0 (60059.0, 60536.0)	4701.0 (4587.0, 4813.0)	34013.0 (33789.0, 34210.0)	386.0 (360.0, 420.0)	613.0 (574.0, 662.0)	0.985 (0.983, 0.986)

In the table each row represents a unique scenario considered by the model. These scenarios are defined by a unique combination of flight time range, and isolation routine (displayed in the left two columns in hours). The first describes how flight times are sampled in that scenario (assumed to always been uniformly distributed between two given limits). We use the ranges [3,5] for short-haul, [7,9] for medium-haul, and [11,13] for long haul flights.

The second column describes the double testing procedure being deployed. The first value is the time (in hours, after arrival into UK) that the second test is administered. Then, the second value is the time (in hours, after arrival into the UK) that travellers shall be released from self-isolation into the general UK population (given that they test negative to both tests and display no symptoms). Note that in the top three rows, this column holds the values (None, None). This represents our base case (for each flight distribution), where no additional isolation or testing is administered, outside of the testing done on arrival.

The remaining values are as described above, presented with their (5, 95)-percentile ranges.

Conclusion

The effectiveness of implementing double screening is largely dependent, by way of positive correlation, on the length of time incoming travellers must spend in self isolation. As the length of time spent in isolation increases, we can also see a sizeable decrease in the number of infected travellers that are detected after the second test is administered. It does appear however, that this correlation seems to plateau after the (192,240) isolation regimen.

There also appears to be a correlation, all be it much less dramatic, between flight time and success rate. Our results indicate that, when no additional screening is being implemented (i.e. in the base case), we may expect to detect an additional 9% of infected travellers that have travelled via long-haul flights, in comparison to infected travellers that have flown to the UK on a short haul flight. This effect becomes lessened however, as more stringent double testing methods are implemented. For example, this effect is reduced to an increase of 0.9% when then (192, 240) isolation routine is being enforced.

Appendix 1 – Model description, assumptions and parameterisation

Model description

To assess the effectiveness of double screening, we have produced a model (which we explore in more detail presently) to calculate the probability that travellers (who have been infected sometime prior to boarding their flight) would be detectable to testing after arriving into the UK and undergoing some period of travel and self-isolation. These probabilities are derived using Monte Carlo simulations which incorporate known aspects of COVID-19's epidemiology.

Explicitly, the model simulates an individual being infected at some random time, $-t_0$ (which we assume to have occurred at any point during the 14 days prior to the traveller's flight; this time is sampled from a uniform distribution across this period), before attempting to board a flight to the UK at time t=0. Each traveller is also randomly determined to be either "symptomatic" or "asymptomatic". This then determines how that traveller's disease progresses. If the traveller is "symptomatic", they are then assigned an incubation time, T_{inc} (randomly sampled according to the distribution described below) after which they will display observable symptoms, as well as being detectable by testing. Otherwise, when a traveller is "asymptomatic", they are given a "detectable-from" time (also represented by T_{inc} in the forthcoming; sampled according to the same distribution described below), after which asymptomatic travellers are detectable only by testing. This mimics the observed behaviour that, after some period, asymptomatic infections become detectable to testing.

After initialisation, each simulated traveller then attempts to board their flight to the UK, where the model proceeds as described:

- If $T_{inc} < t_0$, we assume the infected person has become symptomatic/detectable prior to boarding their flight and are thus detectable. We then assume that, either by means of exit screening or the traveller being too ill to fly, this traveller does not make it onto their flight. These are then recorded as "*non-fliers*".

- If $T_{inc} > t_0$, the traveller proceeds to board their flight that will take T_{flight} hours, where for each traveller, depending on the current scenario, this has sampled from a uniform distribution on the ranges [3,5], [7,9] or [11,13].
- If $T_{inc} < t_0 + T_{flight}$, we assume the infected person has become symptomatic/detectable during their flight. Thus, they will be detected by the testing that is administered to all incoming travellers on arrival. If detected, infected travellers are then removed from the model, being recorded as "*detected, arrival*"
- If we are considering the base case, the model run for each individual stops here. After simulating all our travellers, we assess the base case by taking the ratio of those detected by initial screening, to those who successfully boarded their flight to the UK. This then gives the ratio of travellers detected by this method, and thus an approximation to the probability of detection by the base scenario, given that infected travellers boarded their flight.
- Otherwise, if $T_{inc} > t_0 + T_{flight}$, then the traveller enters self-isolation for a period of T_{iso} hours (where this value is pre-determined by the scenario being considered), after which the traveller must undergo a second test
- If $T_{inc} < t_0 + T_{flight} + T_{iso}$, then the traveller is deemed to have become symptomatic/detectable while in self-isolation. The traveller will then be detected by this second test, being recorded as a "*detected, iso.*"
- If $T_{inc} > t_0 + T_{flight} + T_{iso}$, then the traveller is still undetectable. This traveller however, must still stay in self isolation for another n days (also determined by the current scenario), which emulates the period required for a set of testing results to be returned
- If, however the traveller is marked as "symptomatic" and $T_{inc} < t_0 + T_{flight} + T_{iso} + n$, then the traveller is deemed to have become symptomatic while waiting for results. As symptoms are deemed to always be observable, this person is subsequently stopped from being released into the wider population, recorded as a "*detected, sympto.*"
- In contrast, if the traveller is marked as "asymptomatic" and $T_{inc} < t_0 + T_{flight} + T_{iso} + n$, then as there is no way to know this person is infected at this point (without administering another test), the traveller is released into the general population, recorded as an "*undetected"*
- Else, we have that $T_{inc} > t_0 + T_{flight} + T_{iso} + n$, and the infected traveller has made it through the entire process without being detectable. They are therefore also recorded as an "*undetected*"
- This process is then repeated thousands of times for each scenario, recording the success rate (found by taking the ratio of infected travellers that have been detected during the double testing process, over the total number of infected travellers that successfully travelled to the UK) for each combination of self-isolation period and flight range. Note that this success rate is also an approximation to the probability that infected travellers would be detected by the implemented double testing procedure implemented in each scenario, given that they successfully boarded their flight to the UK.

This process is then repeated for each scenario 100 times (each time simulating 100,000 initially infected travellers attempting to travel to the UK), which are then used to obtain the average success rate (across 10,000,000 simulated individuals), as well as confidence intervals (which are also reported alongside the results presented above).

It is important to bear in mind that this model only considers infected prisoners (every simulated individual is assumed to have already been infected by the time they begin their journey), and thus no conclusions can be made regarding what proportion of the total incoming travellers one can

expected to be infected. This will presumably be somewhat dependent on levels of prevalence in the country of travel origin, although will likely be highly sensitive to other factors. This consideration lays way outside the scope of this work and shall therefore not be mentioned further.

Modelling assumptions

In the following work, we have assumed the following:

- All simulated travellers have, at some point prior to boarding their flight, been infected with COVID-19
- For each traveller, the time of infection is sampled according to uniform distribution over the 14 days prior to boarding their flight
- This model does not consider disease recovery
- All travellers will, at some finite time, become detectable by testing
- The testing method deployed is *perfect*: detecting 100% of symptomatic individuals, and 100% of the "detectable" asymptomatics
- 50% of infected travellers will go on to become asymptomatic; this is randomly assigned to simulated individuals. We therefore refer to "will-be-symptomatic" travellers as having an incubation period, and "will-be-asymptomatic" travellers as having a "time until detectable"
- The distribution of "time until detectable" for asymptomatic travellers is the same as incubation period distribution for symptomatic travellers
- Travellers are only formally tested upon entry into the UK, and at their second testing time
- If a "will-be-symptomatic" traveller becomes symptomatic between the time of their second test and their time of release (out of self-isolation), they shall be detected prior to this release
- Whereas asymptomatic travellers who become detectable between their second test and their release, shall proceed to exit into the general population (as they demonstrate no observable signs to being infected)
- Any travellers that become symptomatic/detectable prior to boarding their flight, do not fly (whether for reasons of exit screening or by being too ill)
- All arriving travellers undergoing the double testing process
- The flight time for each traveller is randomly sampled from a uniform distribution between a given range
- There is no disease transmission at any stage of the model
- Simulated individuals proceed through the model one at a time, and thus enter self-isolation alone

Model Parameters

Our model only uses one externally sourced parameter. This is the incubation period distribution for COVID-19. The distribution used in this model (from which each simulated traveller's incubation period was sampled) was a log-normal distribution, with parameters $\mu = 1.6112$, $\sigma = 0.47238$. This has been taken from the paper *Incubation Period and Other Epidemiological Characteristics of 2019 Novel Coronavirus Infections with Right Truncation: A Statistical Analysis of Publicly Available Case Data,* by N. Linton et al.