



Institute for
New Economic Thinking
AT THE OXFORD MARTIN SCHOOL

A Workable Strategy for Covid-19 Testing: Stratified Periodic Testing rather than Universal Random Testing

Matthew Cleevely, Daniel Susskind, David Vines,
Louis Vines and Sam Wills

15th April 2020

INET Oxford Working Paper No. 2020-08



A Workable Strategy for Covid-19 Testing:

Stratified Periodic Testing rather than Universal Random Testing*

Matthew Cleevely, Daniel Susskind, David Vines, Louis Vines, and Sam Wills**

15 April 2020

Abstract

This paper argues that daily ‘universal random testing’, as recently proposed by Paul Romer, is not likely to be an effective tool for reducing the spread of Covid-19 and resuming economic activity. We find that more than 21% of the population would need to be tested every day to reduce the Covid-19 reproduction rate (R') to 0.75, as opposed to 7% as argued by Romer. We show this using a corrected method for calculating the impact of an infectious person on others, when testing and isolation takes place. Our calculation allows for asymptomatic cases. Instead we propose ‘stratified periodic testing’ as an alternative strategy.

1 Introduction and Summary

Governments around the world are looking for a testing strategy for Covid-19. In a talk given on April 3, Paul Romer proposed population-wide testing and isolation – what we call “universal random testing”.¹ He made the important points that the economic benefit of a speedier recovery would be measured in \$trillions and this would easily justify spending \$billions on testing, and that there is no necessity for such testing to be highly accurate. Unfortunately, as we will show, the model used in that talk is flawed. Our corrected model suggests that Romer’s proposal is unlikely to work in practice. Instead, we show that targeted testing of particular groups, what we call “stratified periodic testing”, and the subsequent isolation of those who test positive, would be a more effective tool in reducing the spread of Covid-19.

* We are especially grateful to David Cleevely, CBE, FEng, Chair of the COVID Positive Response Committee, Royal Academy of Engineering, and Frank Kelly, CBE, FRS, Emeritus Professor of the Mathematics of Systems in the Statistical Laboratory and Fellow of Christ’s College, University of Cambridge, for their editorial contributions to the paper. In addition, we have received many helpful comments and suggestions from Eric Beinhocker, Adam Bennett, Daniela Massiceti, Bob Rowthorn and Stephen Wright.

** Affiliations:

Matthew Cleevely, MEng MPhil MIET, founder of global online booking system 10to8 Ltd.:

matthew@cleevly.com

Daniel Susskind, Fellow in Economics at Balliol College, University of Oxford daniel.susskind@balliol.ox.ac.uk

Louis Vines; data scientist; louisgabrielvines@gmail.com

David Vines; Emeritus Professor of Economics and Emeritus Fellow of Balliol College, and Senior Fellow in the Institute for New Economic Thinking (INET), University of Oxford, and CEPR. david.vines@economics.ox.ac.uk

[Sam Wills](mailto:Sam.Wills@sydney.edu.au), External Research Associate, School of Economics, University of Sydney;

samuel.wills@sydney.edu.au.

¹ See https://bcf.princeton.edu/event-directory/covid19_04/.

Romer uses a simple model to show how random testing of 7 percent of the population per day for evidence of infection (using ‘antigen’ tests) would be sufficient to halt the pandemic. While already a big ask, we argue that this rests on a mistaken calculation, and that a corrected model implies that the required proportion of daily testing would be more than 21 percent of the population – an impossible task. This result indicates that attempting to test at random the entire population would be a waste of testing resources.

Instead we recommend using periodic testing on smaller stratified specific groups. These would include healthcare staff and other key workers, and those at high risk of creating cross-infection (See section 5 below). Like Romer, we believe that high frequency testing of such groups would be manageable. But, unlike Romer, we also believe that only targeted testing would be feasible, and that such testing – on its own - would be a means of stopping the virus spread in the whole population.

So long as testing remains a scarce and relatively expensive resource, we argue that testing of the general population should be reserved for immunity tests (antibody tests), which would allow those that have been infected to get back to work.

The paper is organised as follows. In sections 2 and 3 we use our corrected model to calculate the proportion of daily testing that would be required to halt the spread of the virus. We also use our model to suggest that Romer’s strategy is best thought of as a risky “throttle” strategy. In Section 4 we examine the robustness of our conclusions.

Section 5 sets out what we think is a workable testing strategy: the periodic testing of stratified groups. We argue that a simple method is available for estimating the testing rate required to stop the spread of the virus in each of the targeted groups, and point to Appendix 1 where this method is fully explained. In addition, we argue that testing should be periodic rather than random, since this will enable testing resources to be used more economically: it can increase the efficiency of testing by perhaps 25%. We go on to present some practical suggestions, outlining how our strategy might be implemented.

Section 6 presents our conclusions.

Appendix 1 explains in detail the simple method which we propose for estimating the testing rate required to stop the spread of the virus in the targeted groups. In Appendix 2 we explain why testing periodically, rather than randomly, can make a very big difference to the efficiency of testing. In Appendix 3 we set out how and why we think that Romer made his error.

We should make it clear immediately that our paper is not intended as a criticism. Romer’s focus on the importance of testing, his lecture of 3 April, and his simulations of the spread of the epidemic which we discuss below are all immensely valuable.

2 The basic ideas

The reproduction rate R in an epidemic is the expected number of cases directly generated by any one infectious case. The *basic* reproduction rate, R_0 , is the initial value of R when all individuals are susceptible to infection, and no suppression policies have been applied. For Covid-19, Romer takes $R_0 = 2.5$ and so will we.

We know that an epidemic can only be controlled if the value of R is brought below 1. When that happens, each infected person infects less than one new person, and the epidemic will die out. If that does not happen, more and more people become infected until – in the end – “herd immunity” is achieved.

Romer wants to get the *effective* reproduction rate R' (R prime) down to $R' = 0.75$ from $R_0 = 2.5$, by randomly testing a fraction of the entire population each day and then isolating those who are found to be positive. In Romer’s analysis, which we follow, the effective reproduction rate R' is the product of the basic reproductive rate, R_0 , and the fraction of the infectious population that is not isolated. R' is below R_0 because a fraction of the population is tested each day and those found to be infectious are isolated.

Let ϕ be the proportion of the infectious population which is isolated. Then Romer writes Equation (1):

$$(1) \quad R' = (1 - \phi) R_0.$$

For $R' = 0.75$, Equation (1) implies that that $\phi = 0.7$, *i.e.* that 70 percent of the infectious population is isolated, and that sufficient tests and isolation are carried out to make this possible.

Drawing on his calculations, Romer suggests that this can be done by randomly testing 7% of the population each day. He believes that this would achieve $R' = 0.75$. With a population of 300 million in the US testing on this scale would require about 20 million tests a day. In the UK with a 60 million population, this would require about 8.5 million tests a day. Romer proposes the immediate allocation of \$100bn in the US to make such an outcome possible.

In Appendix 3 we identify the error in Romer’s calculation of the relationship between the proportion of the population who must be tested, and ϕ , the proportion of the infectious population that is isolated. We also explain how we think that Romer came to make his mistake.

We now set out our calculation which shows that, if random testing followed by isolation were adopted, more than 21% of the population would need to be tested each day. This, we show, is what would be necessary to get to a position in which 70 percent of the infectious population were isolated (*i.e.* $\phi = 0.7$), so that $R' = 0.75$. To do this would require everyone in the population to be tested about every five days.

3 Our approach

3.1 Calculating the required testing rate for $R = 0.75$

We assume, like Romer, that the whole population is randomly tested. We let t be the proportion of the population tested each day, *i.e.* the probability, for each person, of being tested each day and then isolated. Romer allows for false negatives in his tests. He lets n be the proportion of false negatives and assumes that this proportion is 0.3. We follow him in assuming such a high number.²

We assume that 14 is the number of days for which an infected person is infectious. This number 14 is familiar in the analysis of Covid-19 as the number of days after which the person is either dead or – much more likely – recovered, but no longer infectious. Of course, this number may not be the best one to use. We use 14 here mainly because it is the number used by Romer as the number of days that each person who tests positive is placed in isolation. We discuss this issue further in Section 4 below, and discuss Romer’s procedure in Appendix 3.³

As a preliminary, let us consider the impact of an infectious person on others. We assume that if $R_0 = 2.5$, and if there were no testing which led to the isolation of infected people, then an infectious person would infect 2.5 people in total, or $2.5/14$ persons per day for 14 days. That is what we take $R_0 = 2.5$ to actually mean. Note that we discount any idea of a person being more or less infectious during the 14-day period. See a brief discussion of this point in Section 4 below.

We now consider the impact of an infectious person on others when he or she has a probability t of being tested each day, and so of being placed in isolation immediately if the test is positive.⁴ For clarity, it is helpful to think about x , where $x = t(1-n)$, where n is the number of tests which show false negatives, which we will take to have the value of 0.3, as specified by Romer. The variable x shows the probability that, on any day on which this person is infectious, he or she is isolated. This means that $(1 - x)$ is the probability that this person will not be isolated on the next day and so will infect people on that next day.

It is important to self-isolate those who show symptoms. Once this happens, only asymptomatic people and those who have chosen not to self-isolate, for example with mild

² There is a good reason for this. Massive testing – even of the amount which we contemplate - may well make it impossible to ensure that tests are accurate. Romer rightly argues that, whilst accurate tests are absolutely necessary for clinical reasons when treating an individual person, much more rough-and-ready testing is satisfactory if the purpose of this testing is epidemiological control.

³ 7 days has been the standard advice for isolation or 14 days if more than one in a household, however recent data shows that the infectious period may last much longer. A recent detailed study of repeatedly tested individuals in Taiwan found a long tail for infectiousness. For further discussion of this point See Section 5 below. See also <https://focustaiwan.tw/society/202003260015>

⁴ This testing regime assumes that on the day you are tested, if tested positive and isolated, you cannot infect someone else. This simplifying assumption can be thought of in practical terms as a rapid ‘early morning’ test. In the section below on robustness we investigate the effect of a much more cautious assumption supposing that the infected individual remains infectious for the entire day that they are tested, with corresponding increase in required testing rate t for any given R' .

or mistaken symptoms, would be spreading the disease, and would be included in those who are tested. Romer discusses this issue in his lecture, but did not allow for it in his calculations.

The proportion of asymptomatic patients isn't really known, and estimates vary wildly. Very inexact data from Iceland suggest all infectious patients are asymptomatic for the first five days and after that only about half become symptomatic.⁵ The WHO suggests that 80% of cases are asymptomatic or mild.⁶

We construct our analysis as follows. We want to find the value of x that would give Romer's desired value for R' of 0.75. We will then work out the required probability of testing t , given that we take as given Romer's assumption that the proportion of false negative tests, n , is equal to 0.3.

If there were *no* self-isolation of those who became symptomatic, then the probability of an infectious person remaining undetected on day k is $(1-x)^k$. Therefore, in expectation, an infected person would infect $R_0(1-x)/14$ persons on day 1 of their infection, $R_0 (1-x)^2/14$ persons on day 2, and so on, up to $R_0 (1-x)^{14} /14$ persons on day 14.

Let us allow for self-isolation of those who are symptomatic by supposing that all of those infected are asymptomatic for five days (and are subject to random testing during that time) and that, after 5 days, a proportion α display symptoms and self-isolate, so that only a proportion $(1 - \alpha)$ go on being tested from day 6 onwards.

We can thus write our following key equation:

$$(3) \quad R' = R_0 [(1-x) + (1-x)^2 + \dots + (1-x)^5 + (1-\alpha)\{(1-x)^6 + (1-x)^7 \dots + (1-x)^{14}\}]/14.$$

Romer's ambition is for each infectious person to only infect 0.75 other people. So we seek to find x by setting the total number of infections caused by this person to 0.75. Recalling that $R_0 = 2.5$, we then can solve for x from Equation (4):

$$(4) \quad R_0 [(1-x) + (1-x)^2 + \dots + (1-x)^5 + (1-\alpha)\{(1-x)^6 + (1-x)^7 \dots + (1-x)^{14}\}]/14 = 0.75.$$

The solution to this equation can be obtained numerically for various values of α . When $\alpha = 0.5$, as roughly observed in Iceland, $x \approx 0.146$. This has the following meaning. Averaging over the first 5 days there is a probability of about 55% for any person who is infectious but not yet isolated, of being tested and immediately isolated. From then on, half of any such infectious people are assumed to display symptoms and so to self-isolate. So, during the remaining 9 days that a person is infected the weighted impact of testing is only about half that in the first 5 days.

But $t = x/(1-n)$. And n , the proportion of false negatives, is equal to 0.3. So the required probability of testing on each day, t , is given by $t = .146/0.7 \approx 0.209$. That is to say, the

⁵ See <https://edition.cnn.com/2020/04/01/europe/iceland-testing-coronavirus-intl/index.html>

⁶ See <https://www.medrxiv.org/content/10.1101/2020.02.20.20025866v2>

proportion of people tested on any day must be as high as 21% in order to achieve the required 15% discovery rate x .

Recall from Equation (1) that $R' = (1 - \phi)R_0$, where ϕ is the proportion of the infectious population which is isolated. With $R_0 = 2.5$ and $R' = 0.75$, this means that $\phi = 0.7$. Thus, over the 14 days in which a person is infectious, this person will, on average, be in isolation for 70 percent of the time. We have shown that to achieve such a very striking outcome, the probability of an infected person who is not yet isolated being tested on any day must be at least as high as 21%. With random, economy-wide testing, this means, in effect, that everybody in the economy has to be tested about every five days.

3.2 Identifying the testing threshold

The dotted line in Figure 1 plots R' as a function of the proportion of the population tested. As we have already seen, when half of those who are infected self-isolate from day 6 onwards, *i.e.* $\alpha = 0.5$, t needs to be equal to about 21% to get R' down to get 0.75. Figure 1 also enables us to identify the threshold testing rates, t^* that reduce R' to exactly 1, and so just stop the epidemic from exploding. For $\alpha = 0.5$, this threshold testing rate is $t^* = 13\%$. To achieve this, everyone would need to be tested, on average, every eight days. This is still way above Romer's proposed testing rate of 7%.

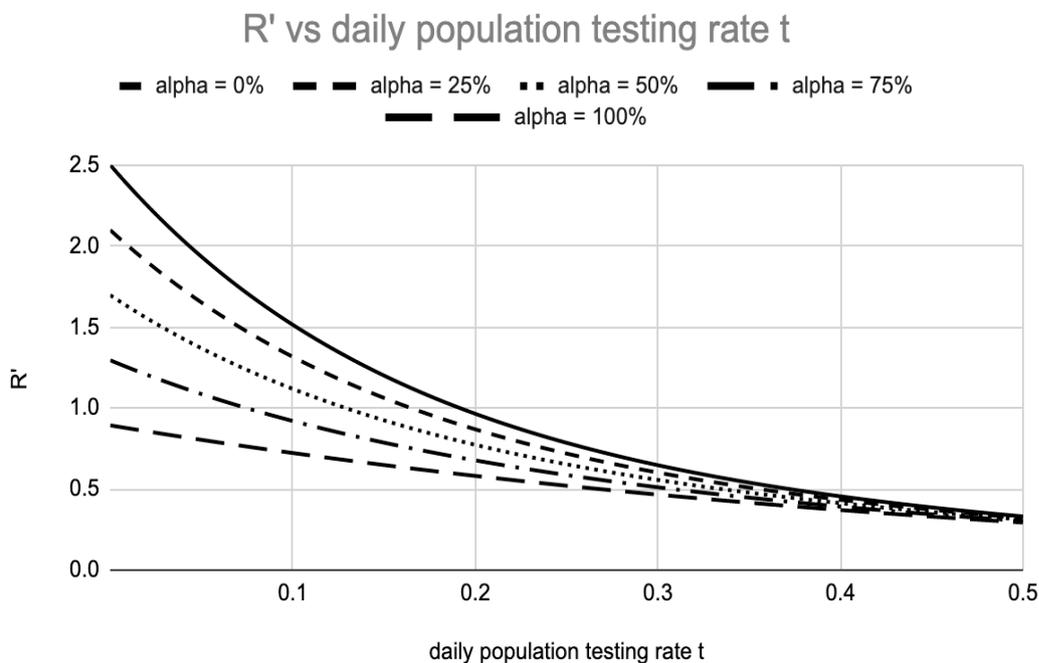


Figure 1

In addition, Figure 1 displays the sensitivity of the results to various values of α . It shows that, at one extreme, when all cases are symptomatic after 5 days and self-isolate (*i.e.* when $\alpha = 1.00$), the situation seems just about manageable: the epidemic would stop exploding even without any testing. Nevertheless, the required value of t which would ensure that $R' = 0.75$, is about 8.2%. This percentage is actually a little *above* that proposed by Romer, even

in this extreme case. That is because infected asymptomatic people do a lot of damage in the first five days! At the other extreme, with $\alpha = 0.00$, the situation is much worse: the probability of testing per day, t^* , which is required to stop the epidemic exploding is now nearly 20 percent (19.1%) and the value of t required to ensure that $R' = 0.75$ is about 26 percent (26.1%).

Data for the extent to which infectious people become symptomatic is extremely unreliable and furthermore the range of possibilities seems very wide. As a result, the outcomes depicted in Figure 1 suggest that, on the balance of probabilities, Romer's strategy of universal random testing would be unworkable.

3.3 Romer's strategy is really a risky "throttle" strategy

Returning to the share of asymptomatic cases observed in Iceland, (*i.e.* the case when $\alpha = 0.5$), Figure 1 shows that by setting a testing rate of 7% of the population, the epidemic remains explosive with R' equal to about 1.3. In this situation, infection would spread rapidly in the earlier stages, since the testing rate is not high enough to slow the spread in a controlled manner; all testing would do is slow down the explosion. Ultimately, once the contagion reaches a certain size, the effect of testing, together with the fact that more and more people have had the disease and so are immune, will begin to slow the spread. The proportion of those infected will tend towards a constant level at which there is "herd immunity". It can be shown, for any R , that this proportion is given by $(1 - 1/R)$.

Without testing, with a value of R_0 of 2.5, the herd immunity proportion is 60%.

Romer's strategy, with a massive amount of testing, would reduce this level to about 23%. But in the earlier stages of infection, when there is no immunity, the disease could still spread rapidly as the testing rate would not be high enough to slow the spread in a controlled manner. Romer's testing strategy would indeed slow the spread. And it would reduce the level of herd immunity, to which the population is tending in the long run. But it would not control the initial explosion.

What Romer is really proposing is a random testing strategy which could be used as a throttle to control an inevitable spread, and that at 7% this throttle strategy might be 'good enough'. But such a control mechanism does not stop very large numbers of people being infected, it merely "flattens the curve". Of course, the hope is that something else intervenes, like a vaccine.

Romer has posted a very detailed and helpful model of the spread of the epidemic in this manner, one which avoids the problems of the model put forward in his April 3 talk. See <https://paulromer.net/covid-sim-part1/> and <https://paulromer.net/covid-sim-part2/>.

In Romer's simulations of this model, the spread is very rapid: there is a peak of infections of between 3% and 18% of the population. This would overwhelm any national health system

since even 3% of the population is an enormous number. There is also a chance that 20% of the population would be infected at the same time. That would be a national calamity.

Romer shows in his simulation model that using such a strategy over the course of 500 days might result in about 30% of the population contracting Covid-19. Furthermore his strategy is risky. Peak levels of infection might rise out of control. Even if testing rates were then increased, lags in responses would mean that the spread of infection would only be gradually reduced. Meanwhile, the virus might go on spreading towards the herd immunity level.

4 The robustness of our conclusions

Of course, there are many changes to our assumptions which could modify our calculations.

In particular, there would be significant *reductions* in required testing rates if whole households were to be self-isolated if anyone in the household tested positive. If, for example, only one person in a household were tested at any time, and households consisted on average of two people, then each positive test would remove two people into isolation. That is – testing would become more effective.

On the other hand, our calculations have deliberately assumed a very speedy testing strategy: we have supposed that a test done on any day which finds the person to be infectious causes that person to immediately self-isolate even on that day; an extreme assumption. We have redone the calculations using the more cautious assumption that a positive test result for a test performed on any day does not lead to the person isolating until the next day. The relevant equation now becomes:

$$(5) \quad R' = R_0[1 + (1 - x) + (1 - x)^2 + \dots + (1 - x)^4 + (1 - \alpha)\{(1 - x)^5 + (1 - x)^6 \dots + (1 - x)^{13}\}]/14$$

These results are *much* worse than those described in section 3. With $\alpha = 0.5$, the critical testing rate, t^* , is now 16% percent, and the testing rate required to get R' down to 0.75 is now as high as 27% percent.

Furthermore, we have been assuming, like Romer, that there is uniform contagion throughout the 14-day period. But the medical data shows an asymptomatic infectious period followed by a hump of maximum infectivity as symptoms develop and a tail as symptoms resolve.

Our conclusions from this investigation are as follows. For the random universal testing proposed by Romer to be workable (involving testing, say, less than 10% of the population per day) policy-makers would require confidence that:

- i) nearly all infected patients are symptomatic and self-isolate, reducing the burden on testing after the incubation period,

- ii) the tests are sufficiently effective, and complied with, that they capture more than 70% of infected cases (and ideally close to 100%),
- iii) testing is conducted quickly, early in the morning, and people are isolated on the day of the test, and
- iv) whole households are isolated when any member is infected.

Unfortunately, we do not feel that all these conditions can be met given our current state of knowledge about the virus, particularly i). Thus we think that it is right to be very cautious before agreeing with Romer that whole-population random testing would be helpful.

5 A workable strategy of stratified periodic testing⁷

5.1 Stratified testing

We argue that testing should be carried out at different frequency for different stratified groups, based on occupation, geography, or other factors. Testing at rates above 25% per day could be done for carefully selected groups for which R_0 was high. This would enable greater rates of isolation in these groups, lowering their R' and helping to prevent the epidemic spreading where it mattered most.⁸ This appears to be a much lower-risk strategy to contain the spread of infection, and could be done with cheap tests, even if they were somewhat inaccurate.

Consider doctors, for example, who have very frequent contact with infected patients. The value of R_0 will be very high for doctors. The above kind of calculation could be carried out to ensure that the effective reproduction rate, R' , for doctors in hospitals can be brought well below 1. It appears likely that such calculations will show that doctors need to be tested every day. This is something which Romer already suggested in his talk.

Testing health workers is of course obvious and accepted. Other groups needing to be tested include those at high risk of cross infection (such as the elderly in care homes), and those with high exposure who are both vulnerable themselves and likely to cross infect (such as bus drivers, grocery store workers, and police).

The kinds of calculation described in Appendix 1 could easily be carried out for structured samples in different locations, so as to respond to different patterns of infectivity. Some parts of the population will have low R_0 and need significantly less frequent testing. Testing should be concentrated in areas in which R_0 is high.

⁷ Some of what follows comes from suggestions made to us by Eric Beinhocker, for which we are very grateful.

⁸ It might even enable the general lockdown to be eased, so that other lower-risk groups could keep working and not need to be isolated.

5.2 Clarifying how our strategy differs from that of Romer

We hope that Romer would agree with what we have just written above. Indeed, in the version of his plan that he set out on Twitter, Romer has himself provided useful suggestions about who might have priority as tests are rolled out.⁹ But from here on we part company.

Romer goes on to suggest that, once tests have been rolled out for these most important groups, there be a further vast expansion of testing, enabling mass random testing *for the whole population*, in order to get R down *for the whole population*. The version of his plan on Twitter makes this very clear. It concludes as follows:

“When you strip away all the noise and nonsense, note that once we cover essential workers, it’s easy to test everyone in the US once every two weeks. Just do it. Isolate anyone who tests positive. Check your math. Surprise, $R_0 < 1$. Pandemic is on glide path to 0. No new outbreaks. No need for any more shutdowns.”¹⁰

Instead, we argue that testing must be focused on particular groups. This is because our findings, discussed in Section 3 above, show that a mass testing plan would still leave R' significantly above 1 unless it was carried out incredibly frequently.

Nevertheless there will still need to be random testing of groups in the population, and some random testing of the whole population. But this testing would be *for informational* purposes only, and would only involve testing very small samples of those involved.

Such informational testing will be needed for two reasons. First, random testing of small samples from particular groups will be necessary to track the groups in which R_0 is high, where there is a potential for a high rate of spread. Once identified, these particular groups will then need very frequent testing of everyone in the group, for the reasons which we have been discussing in this paper. Second, testing of small samples of wider groups in the whole population will be needed to find where prevalence is already high – so that spread has already become rapid. Such groups will also need very frequent testing of everyone in the group. Doing all this will also help governments to track spread and to determine where hotspots are flaring up. Such information will help governments to work out how to selectively tighten, or loosen, containment measures when needed.

But the accuracy of this testing for informational reasons will be determined by the sample size, rather than population size. The samples required for these informational purposes will be *very small* relative to the size of the whole population.

⁹ See <https://threadreaderapp.com/thread/1248712889705410560.html>

¹⁰ Romer uses R_0 to stand for what we call R' See again <https://threadreaderapp.com/thread/1248712889705410560.html>

5.3 A simple way of calculating threshold testing rates

It would be good to find a simple way of calculating the testing threshold for any group, *i.e.* the value which t must exceed in order to ensure that R is less than 1, and so prevent the epidemic from exploding in that group. In fact, we can do this by using a simple approximation which ignores the dynamics of the infection process, thereby producing an equation which is easy to solve. By ignoring these dynamics, we do not have to solve a complex equation like Equation (3) which sums a number of effects in a non-linear way, over many time periods. This is likely to be helpful.

In Appendix 1, we set out our simple approximate method for calculating the threshold testing rates t^* , for different groups, based on their initial values of R_0 . We show that the approximation is relatively accurate, even although it rests on two simplifying assumptions. Because our proposed calculation is so simple it appears to be useful. Furthermore, for large values of α , it gives results which actually *overstate* the required amount of testing. Such a conservative approach seems desirable, in the light of the current uncertainty about the true value of α in populations, and the strong impact that this value will have on the infection rate.

5.4 Periodic tests rather than random tests

Once the frequency of testing has been decided and testing kits are available, testing can begin for everyone in the identified groups. But it is important that this testing be done periodically for each person, rather there being a random choice of those who are to be tested in each time period.

The rationale underlying periodic testing can be explained by the “waiting-time paradox”.¹¹ Random testing, say of 25 percent of a group each day, wastes many resources. This is because every day some of those tested will have been tested the previous day, whilst others who have not been tested for a long time will, nevertheless, still not be tested and so will possibly continue infecting people. By contrast, periodic testing of 25 percent of the group means that, on days one to four, a different quarter of the group will be tested each day, and that on day 5 the first quarter of the group will be tested again, *etc.* It is clear that this means each person tested will have been tested exactly four days previously, removing the problem that some tests are being wasted and that some other tests will be postponed for too long.

In Appendix 2, we provide a simple account of this issue, and show how important it is likely to be. We show that with high testing rates, periodic testing beats random testing by a very significant factor. For example, in the model which we examined in Section 3, in the special case in which there is no self-isolation of symptomatic people (*i.e.* with $\alpha = 0$), our testing

¹¹ The waiting time for a Poisson bus service is twice the waiting time for a periodic bus service with the same rate for a randomly arriving traveller.

rate required for R_0 to be 0.75 was 18.3% with random testing. With periodic testing this rate falls to 13.5%, a 26% reduction.

This is a big improvement at *no* extra cost.

5.5 Running two kinds of tests in parallel

In this paper, we have been discussing antigen testing (i.e. testing for active infections) as opposed to antibody testing (i.e. testing for those who have had the disease and are immune). A combination of the two kinds of tests might be used in a way that would be effective and realistic if one assumes that capacity of testing for active infection remains constrained, but that one-time antibody tests become widely available. One might then proceed as follows:

- Immediate active infection testing for high-risk groups.
- Self-isolation for anyone else developing symptoms. Isolation would, as currently, last for a minimum of 7 days and the individual would not be tested unless this was medically necessary. But if there were enough tests then one could test people at the *end* of their isolation period to show that they were clear of virus before they were allowed to come out of isolation.
- Widespread home kit antibody testing for anyone to see if they had had the virus - these would be one-off tests that would not need to be repeated.
- A system to track people with immunity who could then circulate freely if they had either a) had a positive antibody test, b) had a positive active infection test more than some specified number of days ago, or c) had a negative active infection test after their symptoms resolved.

All these measures would be realistic in terms of the required testing resources. Such a testing procedure would involve doing two things at once: the stratified periodic antigen testing which we have been discussing would be designed to damp the spread of the disease in key groups, by catching those in these groups who were infectious but asymptomatic, or pre-symptomatic, or post-symptomatic, and so not self-isolating. At the same time, antibody testing for the entire population would separate out the immune population; passing an antibody test would enable such people to return to work.

6 Conclusions

Romer suggests that by testing 7% of the population every day we can get the effective reproduction number of Covid-19 to around 0.75 and curb the epidemic. Unfortunately, his calculations contain errors. By correcting his method, and using reasonable assumptions about asymptomatic carriers, we believe that at least 21% of the population would need to be tested each day to get the reproduction number well below 1. For obvious reasons, we do not see this as a feasible population-wide strategy.

Instead, we argue that cheap, less reliable, tests will work for narrow groups. Testing rates of over 25% per day in certain groups, such as health workers, could eliminate the spread, enabling other lower risk groups to keep working and out of isolation. Our work has also shown that periodic testing would be a more effective and workable solution than random testing.

Additionally, if groups of high prevalence can be identified (say by geography or occupation) more frequent testing will both ensure that any infected person is isolated more quickly and that more infected people test positive and are isolated, thus maximising the effectiveness of any testing.

Appendix 1 A simple method for calculating minimum testing rates

In this Appendix we set out a simple method for calculating the testing threshold in any group, *i.e.* the value, t^* , which t must exceed in order to prevent the epidemic from exploding in that group. The calculation employs a simple approximation which ignores the dynamics of the infection process. If we ignore the dynamics, we do not have to solve a complex equation like Equation (3) which sums a number of effects in a non-linear way, over many time periods.

Our method of calculation builds on the following insight: for R' to be less than 1 the following must be true: on any given day a person with Covid-19 is more likely to go into isolation than to spread it to someone else. If that is true, then the expected number of infections a person will hand on will be at most less than 1. We can ignore the dynamics of the process and simply solve for the value of t for which this will be the case. We proceed in two steps.

(a) For simplicity, we first examine the extreme case in which none of those who are infected become symptomatic and self-isolate; this corresponds to the case considered in Section 3 in which $\alpha = 0$.

Consider any group of z , as yet unidentified, infectious people. Assuming that this group is a small fraction of the overall population, the number of people who will be infected by this group on any given day is (R_0/d) times z , where d is the number of days that an infectious person remains infectious.

The number of these z people who, on this same day, will go into isolation because they have tested positive will be $t(1-n)$ times z . But there will be additional infectious people who cease to infect others because, although they did not test positive on that day, the period during which they had the disease and were infectious will have come to an end. This happens with probability $1/d$; so there will be $\{[1-t(1-n)]/d\}$ times z such people¹².

¹² This $1/d$ probability is the chance that an infected person becomes non-infectious independently of testing. An intuitive way to think about this is that, in choosing someone at random, there is a $1/d$ chance that that person is on their last day of infection and so will become non-infectious the following day. This is only an

Thus, for R' to be less than 1, we require that:

$$(6) \quad \{t(1-n) + [1-t(1-n)]/d\} > R_0/d.$$

This means that, for this extreme case, the threshold testing rate is given by

$$(7) \quad t^* = (R_0 - 1)/[(d-1)(1-n)]$$

If $R_0 = 2.5$, $d = 14$, and $n = 0.3$, we get $t^* = 16.5\%$. That is, this method says that, to get R less than 1 by randomly testing the whole population, one needs to test at least 17% of the population. That is, the threshold testing rate, t^* , is 17%.

This is a lower value than what we found for t^* using the full dynamic model in Section 3 when $\alpha = 0$. The result there was that $t^* = 19.1\%$. The discrepancy between these two results arises precisely because of the dynamic process of the epidemic: the simple calculation carried out here ignores the fact that, as time passes, testing will remove some of the infected people, so that they are no longer available to be tested on later days; that is what made Equation (3) so complex! But this simple method ignores that fact.¹³ For this reason the result produced using this simple method will always underestimate the required testing rate. This method thus provides a (quick and dirty) lower bound for the true value of t^* . Nevertheless it may provide a useful starting point.

(b) This method can be readily extended to include the more general cases considered in Section 3 in which a proportion of those who are infected become symptomatic after a certain number of days and so self-isolate. Consider here a proportion α who self-isolate after a number of days d_0 out of the total number of days of infection d . We can approximate an adjusted value of α , which we call α' . This is the probability that someone who is infected self isolates on any particular day (independently of any test or of reaching the end of their infectious period), such that at the end of the infectious period the chance the individual has self-isolated is α , namely;

$$(8) \quad \alpha' = \alpha/d$$

We then introduce α' in equation (6) to create a new condition of the threshold testing rate, t , for a population following this self-isolation rule:

$$(9) \quad \{t(1-n) + [1-t(1-n)]/d + (1-t(1-n)) (1-\alpha')/d\} > R_0/d.$$

approximation since it requires that the value of R' resulting from testing is equal to 1. That is because if $R' > 1$ then the virus would be spreading and hence an individual would be less likely to be on their last day of infection; conversely if $R' < 1$ then a higher proportion of the infected population would more likely to be about to end their infectious period.

¹³ It is possible that this is what Romer was effectively assuming in his analysis. See the final paragraph of Appendix 3.

This means that the threshold testing rate is now given by

$$(10) \quad t^* = (R_0 - d\alpha' - 1 - \alpha') / [(d - d\alpha + \alpha' - 1)(1-n)]$$

Suppose, as in the previous case, that $R_0 = 2.5$, $d = 14$, and $n = 0.3$. Then for a population for whom the first 5 days are asymptomatic who then become symptomatic and self-isolate with probability $\alpha=0.5$, Equations (7) and (8) produce an adjusted α' for the population of 0.023 (those self-isolating on any given day). This leads to a correspondingly reduced threshold testing rate of $t^* = 11.0\%$. In other words, this method says that, to get R less than 1 by randomly testing the whole population, one would only need to test 11% of the population.

This is a smaller value than what we found for t in Section 3, in this case with $\alpha = 0.5$, using the full dynamic model. The result there was that $t^* = 13\%$. The discrepancy between these two results arises for the same reason that it did in the case in which $\alpha = 0$: the simple calculation carried out here ignores the fact that, as time passes, testing will remove some of the infected people, so that they are no longer available to be tested on later days. However, we have found that, for larger values of α , the simple model being examined here will actually *overstate* the testing required. That is because we are here making a second simplifying assumption, that there is a 'random' self-isolation process α' each day, such that at 14 days (d) the chance that someone has self-isolated is exactly equal to α . This is as opposed to the detailed model in Section 3 in which, after day 5, there is a step of size α in the chance of someone self-isolating, so that a proportion α self-isolate for a 9 day period with certainty. Nevertheless, even with these two simplifying assumptions, the result obtained here appears to be useful. We say this in the light of the current uncertainty about the true value of α in populations and the strong impact this will have on the infection rate. It seems appropriate to use a calculation method which, when α is high, provides an overly high result for the testing threshold. Acting cautiously in these circumstances seems desirable.

Appendix 2: Periodic testing versus Random Testing¹⁴

Our aim is to show how much more effective periodic testing might make the testing process, when compared with the random testing process discussed in Section 3. As in that case our objective is to get the value of R down from $R_0 = 2.5$ to $R' = 0.75$. For simplicity, in this Appendix, we will deliberately ignore the effects of self-isolation, *i.e.* the results in Section 3 with which we must compare our findings here are those in which $\alpha = 0$. We also suppose in this Appendix, again for simplicity, that tests are conducted perfectly.

¹⁴ We are grateful to Frank Kelly for his help in preparing this Appendix

Our results suggest that periodic resting might be about 25% more effective than random testing, at no extra cost.

As in Section 3, we let R' be the expected number of people that a randomly chosen infected person infects before that person is positively tested (or stops being infective, if sooner). Let r_j be the expected number of individuals infected by an individual on day of his/her infection, for $j=1,2,\dots,d$ where the length of infectivity is d . (This is the same as l in the model in Section 3). Thus $R' = \sum_{j=1}^d r_j$. Now suppose an individual is tested every N days, and that for high risk groups we test very frequently, so that $N < d$. If the time of the infection is random and the individual is *not* infective for the day of the test, then

$$(1) \quad R' = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{i-1} r_j$$

From this we can deduce that

$$(2) \quad R' = \frac{1}{N} \sum_{j=1}^N (N - j + 1) r_j$$

For example, if $r_j = R_0/d$ then¹⁵

$$(3) \quad R' = R_0 \frac{N+1}{2d}$$

We can now compare our findings with those in Section 3, for the case in which $\alpha = 0$. We found there that that to bring R down from $R_0 = 2.5$ to $R' = 0.75$ would – with random testing - require a testing rate of 26 percent. But those results assumed that 30 percent of tests failed. If tests were perfect those results imply a testing rate of 18.3 percent.

We can now use Equation (3) immediately above to solve for the value of N , the number of days between each test, that is required to bring R down from $R_0 = 2.5$ to $R' = 0.75$. This equation shows that the period between testing for each individual would need to be 7.4 days.¹⁶ That is a testing rate of 13.5 percent.

This is a twenty-six percent reduction in the rate of testing required, as compared with the case of random testing. We can see that doing random testing would provide a big improvement at *no* extra cost.

If scarce testing is to be allocated over individuals with different prior probabilities of infection, these formulas can be used to optimize the allocation. They show that there are diminishing returns from making N very small. Suppose that individual k has prior probability p_k of infection and that we wish to choose N_k so as to allocate a given amount

¹⁵ For low-risk individuals we can test much less frequently, so that $N \geq d$. Then we obtain

$$(4) \quad R' = \frac{1}{N} \left(\sum_{i=1}^M \sum_{j=1}^{i-1} r_j + (N - d)R_0 \right)$$

and hence

$$(5) \quad R' = R_0 - \frac{1}{N} \sum_{j=1}^d j r_j$$

For example, if $r_j = R_0/d$ we can solve for the required rate of testing. It is

$$(6) \quad R' = R_0 \left(1 - \frac{d+1}{2N} \right)$$

¹⁶ Of course in reality such a number would need to be rounded up or down to a full number of days.

of testing over a set of individuals so as to maximally reduce R' . Then, using the last displayed formula, the optimal allocation of a given amount of testing should choose $N_k \propto 1/\sqrt{p_k}$. The amount of testing allocated to higher risk individuals is naturally higher, but not proportionately so.

Appendix 3: Romer's analysis of testing

We now explain -why we think there is a mistake in the way in which Romer calculates ϕ , the proportion of the infectious population which is isolated. We then present our attempt to understand how and why he made his error.

A.2.1 Romer's analysis

Romer assumes random testing of the whole population. In his calculations, he lets t be the proportion of the population tested each day, *i.e.* the probability that each person is tested each day. He supposed that $t = 0.07$.

Romer allows for false negatives in tests. He lets n be the proportion of false negatives. Romer assumes that this proportion is 0.3.

Romer lets l be the number of days that each person who tests positive is placed in isolation. He assumes $l = 14$.

Romer then computes ϕ , the proportion of the infectious population which is isolated, as follows. He writes something similar to, but not the same as what we have called Equation 1 in our paper.

$$(1) \quad \phi = t(1 - n)l$$

Just to be clear, this equation here comes directly from the slides which accompanied Romer's talk¹⁷. We have no background on why he wrote down this equation, and we think that it is incorrect. We say this because in our paper above Equation (1) reads $\phi = t(1 - n)d$. This has the variable d on the right-hand side, showing the number of days for which a person remains infectious. By contrast Equation (1) above has l on the right hand side the variable l which is the number of days that each person who tests positive is placed in isolation. The nature of what we think is Romer's error is discussed immediately below.

Since $t = 0.07$; $(1 - n) = 0.7$ and $l = 14$ Romer claims that $\phi = 0.69$. This value of $\phi = 0.69$ would, he says, produce his desired value for R' , since:

$$R' = (1 - \phi) R_0, \text{ or } R' = (1 - 0.69) \times 2.5 \approx (1 - 0.7) \times 2.5 = 0.75.$$

Drawing on these calculations, Romer suggests that there should be testing of 7% of the population each day.

Notice that, although Romer mentioned self-isolation of those who have symptoms in his lecture, there is no allowance for such an action in any of the calculations in his slides.

¹⁷ See minutes 16 to 20 of the Romer talk, and the accompanying slides.

A.2.2 Our Criticism

It is helpful to try to understand how Romer made what we think is an error.

To see most clearly, and simply, why his calculation cannot be right, imagine what would happen if there were to be double the amount of testing proposed by Romer, *i.e.* suppose that $t = 0.14$. Then, using his formula for ϕ we would get $\phi = t(1 - n)l = 0.14 \text{ times } 0.7 \text{ times } 14 = 1.38$; the person would be in isolation for more than all of the period of 14 days! So the equation must be wrong.

How can we understand the inclusion of 'l', the number of days that an infected person is placed in isolation, on the right-hand side of this equation? One possibility is that the inclusion of 'l' is simply a mistake. Romer states that $\phi = t(1-n)l$, but ϕ and $t(1-n)l$ appear to be very different things. Because $t(1-n)$ is equal to the probability that an infectious person is put into isolation on any day, it follows that $t(1-n)l$ is equal to the expected length of isolation any infected person is likely to face, after one round of testing. But ϕ is the fraction of the infected population that are isolated -- which is clearly not the same thing as the expected length of isolation any infected person is likely to face, $t(1-n)l$. So it appears that stating $\phi = t(1-n)l$ is a mistake.

Another possibility is to make a set of assumptions about Romer's set-up that bring the meaning of ϕ and $t(1-n)l$ closer together. For instance, consider the following approach. First, interpret 'l' as the 'number of days that an infected person is infectious', rather than 'is placed in isolation'. Secondly, define 'Z' as the number of infected people not in isolation. Thirdly, imagine there are 'l' periods where, in each period, a fraction $t(1-n)$ of those infected people not in isolation, Z, are removed and put into isolation. And finally, assume that in each period the number of infected people not in isolation, Z, remains the same (*i.e.* the infected who are put into isolation are replaced with newly infected people). Then it follows that, after 'l' periods, $t(1-n)l * Z$ people will be in isolation. Now, $t(1-n)l$ is indeed equal to ϕ , the proportion of the infected population not in isolation who are put into isolation -- but with two very significant caveats. First, it assumes that everyone who will be isolated over the 'l' days is isolated on the first day. And secondly, because Z is constant over time, it follows that ϕ may also be greater than one if t or l is large enough, or n is small enough -- which, as shown before, is exactly the problem with Romer's analysis.¹⁸

¹⁸ Intuitively, the problem here is that you are taking a fraction $t(1-n)$ of the infected population not in isolation, Z, and putting them into isolation in each period, but because l replenishes over time, if you isolate a large enough proportion of Z, $t(1-n)$, enough times, l, then you will end up with more infected people in isolation, $t(1-n)l * Z$, than there are infected people not in isolation, Z.