Gowers's Weblog

Mathematics related discussions

How long should a lockdown-relaxation cycle last?

The adaptive-triggering policy.

On page 12 of a document put out by Imperial College London (https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf), which has been very widely read and commented on, and which has had a significant influence on UK policy concerning the coronavirus, there is a diagram that shows the possible impact of a strategy of alternating between measures that are serious enough to cause the number of cases to decline, and more relaxed measures that allow it to grow again. They call this *adaptive triggering*: when the number of cases needing intensive care reaches a certain level per week, the stronger measures are triggered, and when it declines to some other level (the numbers they give are 100 and 50, respectively), they are lifted.

If such a policy were ever to be enacted, a very important question would be how to optimize the choice of the two triggers. I've tried to work this out, subject to certain simplifying assumptions (and it's important to stress right at the outset that these assumptions are questionable, and therefore that any conclusion I come to should be treated with great caution). This post is to show the calculation I did. It leads to slightly counterintuitive results, so part of my reason for posting it publicly is as a sanity check: I know that if I post it here, then any flaws in my reasoning will be quickly picked up. And the contrapositive of that statement is that if the reasoning survives the harsh scrutiny of a typical reader of this blog, then I can feel fairly confident about it. Of course, it may also be that I have failed to model some aspect of the situation that would make a material difference to the conclusions I draw. I would be very interested in criticisms of that kind too. (Indeed, I make some myself in the post.)

Before I get on to what the model is, I would like to make clear that I am not *advocating* this adaptive-triggering policy. Personally, what I would like to see is something more like what Tomas Pueyo calls The Hammer and the Dance (https://medium.com/@tomaspueyo/coronavirus-the-hammer-and-the-dance-be9337092b56): roughly speaking, you get the cases down to a trickle, and then you stop that trickle turning back into a flood by stamping down hard on local outbreaks using a lot of testing, contact tracing, isolation of potential infected people, etc. (This would need to be combined with other measures such as quarantine for people arriving from more affected countries etc.) But it still seems worth thinking about the adaptive-triggering policy, in case the hammer-and-dance policy doesn't work (which could be for the simple reason that a government decides not to implement it).

A very basic model.

Here was my first attempt at modelling the situation. I make the following assumptions. The numbers S, T, λ, μ are positive constants.

1. Relaxation is triggered when the rate of infection is S.
2. Lockdown (or similar) is triggered when the rate of infection is T.
3. The rate of infection is of the form $Ae^{\lambda t}$ during a relaxation phase.
4. The rate of infection is of the form $Be^{-\mu t}$ during a lockdown phase.
5. The rate of "damage" due to infection is θ times the infection rate.
6. The rate of damage due to lockdown measures is η while those measures are in force.

For the moment I am not concerned with how realistic these assumptions are, but just with what their consequences are. What I would like to do is minimize the average damage by choosing S and T appropriately.

I may as well give away one of the punchlines straight away, since no calculation is needed to explain it. The time it takes for the infection rate to increase from S to T or to decrease from T to S depends only on the ratio T/S. Therefore, if we divide both T and S by 2, we decrease the damage due to the infection and have no effect on the damage due to the lockdown measures. Thus, for any fixed ratio T/S, it is best to make both T and S as small as possible.

This has the counterintuitive consequence that during one of the cycles one would be imposing lockdown measures that were doing far more damage than the damage done by the virus itself. However, I think something like that may actually be correct: unless the triggers are so low that the assumptions of the model completely break down (for example because local containment is, at least for a while, a realistic policy, so national lockdown is pointlessly damaging), there is nothing to be lost, and lives to be gained, by keeping them in the same proportion but decreasing them.

Now let me do the calculation, so that we can think about how to optimize the ratio T/S for a fixed S.

The time taken for the infection rate to increase from S to T is $s = \lambda^{-1}\log(T/S)$, and during that time the number of infections is
$\int_0^s Se^{\lambda t}dt = \lambda^{-1}S(e^{\lambda s} - 1) = \lambda^{-1}(T - S)$.

By symmetry the number of infections during the lockdown phase is $\mu^{-1}(T-S)$ (just run time backwards). So during a time $(\lambda^{-1}+\mu^{-1})\log(T/S)$ the damage done by infections is $\theta(\lambda^{-1}+\mu^{-1})(T-S)$, making the average damage $\theta(T-S)/(\log T-\log S)$. Meanwhile, the average damage done by lockdown measures over the whole cycle is $\mu^{-1}\eta/(\lambda^{-1}+\mu^{-1})=\eta\lambda/(\lambda+\mu)$.

Note that the lockdown damage doesn't depend on S and T: it just depends on the proportion of time spent in lockdown, which depends only on the ratio between λ and μ. So from the point of view of optimizing S and T, we can simply forget about the damage caused by the lockdown measures.

Returning, therefore, to the term $\theta(T-S)/(\log T-\log S)$, let us say that $T=(1+\alpha)S$. Then the term simplifies to $\theta\alpha S/\log(1+\alpha)$. This increases with α, which leads to a second counterintuitive conclusion, which is that for fixed S, α should be as close as possible to 0. So if, for example, $\lambda=2\mu$, which tells us that the lockdown phases have to be twice as long as the relaxation phases, then it would be better to have cycles of two days of lockdown and one of relaxation than cycles of six weeks of lockdown and three weeks of relaxation.

Can this be correct? It seems as though with very short cycles the lockdowns wouldn't work, because for one day in three people would be out there infecting others. I haven't yet got my head round this, but I think what has gone wrong is that the model of exponential growth followed instantly by exponential decay is too great a simplification of what actually happens. Indeed, data seem to suggest a curve that rounds off at the top rather than switching suddenly from one exponential to another — see for example Chart 9 from the Tomas Pueyo article linked to above. But I think it is correct to conclude that the length of a cycle should be at most of a similar order of magnitude to the "turnaround time" from exponential growth to exponential decay. That is, one should make the cycles as short as possible provided that they are on a timescale that is long enough for the assumption of exponential growth followed by exponential decay to be reasonably accurate.

What if we allow a cycle with more than two kinds of phases?

So far I have treated λ and μ and η as parameters that we have no control over at all. But in practice that is not the case. At any one time there is a suite of measures one can take — encouraging frequent handwashing, banning large gatherings, closing schools, encouraging working from home wherever possible, closing pubs, restaurants, theatres and cinemas, enforcing full lockdown — that have different effects on the rate of growth or decline in infection and cause different levels of damage.

It seems worth taking this into account too, especially as there has been a common pattern of introducing more and more measures as the number of cases goes up. That feels like a sensible response — intuitively one would think that the cure should be kept proportionate — but is it?

Let's suppose we have a collection of possible sets of measures $M_1, \dots, M_r$. For ease of writing I shall call them measures rather than sets of measures, but in practice each M_i is not just a single measure but a combination of measures such as the ones listed above. Associated with each measure M_i is a growth rate λ_i (which is positive if the measures are not strong enough to stop the disease growing and negative if they are strong enough to cause it to decay) and a damage rate η_i. Suppose we apply M_i for time t_i. Then during that time the rate of infection will multiply by $e^{\lambda_i t_i}$. So if we do this for each measure, then we will get back to the starting infection rate provided that $\lambda_1 t_1 + \dots + \lambda_r t_r = 0$. (This is possible because some of the λ_i are negative and some are positive.)

There isn't a particularly nice expression for the damage resulting from the disease during one of these cycles, but that does not mean that there is nothing to say. Suppose that the starting rate of infection is S_0 and that the rate after the first i stages of the cycle is S_i. Then $S_i = e^{\lambda_i t_i} S_{i-1}$. Also, by the calculation above, the damage done during the ith stage is $\theta\lambda_i^{-1}(S_i - S_{i-1})$.

In what order should the M_i be applied?

This has an immediate consequence for the order in which the M_i should be applied. Let me consider just the first two stages. The total damage caused by the disease during these two stages is

$$\theta(\lambda_1^{-1}(S_1 - S_0) + \lambda_2^{-1}(S_2 - S_1))$$

$$= \theta S_0(\lambda_1^{-1}(e^{\lambda_1 t_1} - 1) + \lambda_2^{-1} e^{\lambda_1 t_1}(e^{\lambda_2 t_2} - 1)).$$

To make that easier to read, let's forget the θS_0 term (which we're holding constant) and concentrate on the expression

$$\lambda_1^{-1}(e^{\lambda_1 t_1} - 1) + \lambda_2^{-1} e^{\lambda_1 t_1}(e^{\lambda_2 t_2} - 1).$$

If we reorder stages 1 and 2, we can replace this damage by

$$\lambda_2^{-1}(e^{\lambda_2 t_2} - 1) + \lambda_1^{-1} e^{\lambda_2 t_1}(e^{\lambda_1 t_2} - 1).$$

This is an improvement if the second number is smaller than the first. But the first minus the second is equal to

$$(\lambda_2^{-1} - \lambda_1^{-1})(e^{\lambda_1 t_1} - 1)(e^{\lambda_2 t_2} - 1),$$

so the reordering is a good idea if $\lambda_1 > \lambda_2$. This tells us that we should start with smaller λ_i and work up to bigger ones. Of course, since we are applying the measures in a cycle, we cannot ensure that the λ_i form an increasing sequence, but we can say, for example, that if we first apply the measures that allow the disease to spread, and then the ones that get it to decay, then during the relaxation phase we should work from the least relaxed measures to the most relaxed ones (so the growth rate will keep increasing), and during the suppression phase we should start with the strictest measures and work down to the most relaxed ones.

It might seem strange that during the relaxation phase the measures should get gradually more relaxed as the spread worsens. In fact, I think it *is* strange, but I think what that strangeness is telling us is that using several different measures during the relaxation phase is not a sensible thing to do.

Which sets of measures should be chosen?

The optimization problem I get if I try to balance the damage from the disease with the damage caused by the various control measures is fairly horrible, so I am going to simplify it a lot in the following way. The basic principle that there is nothing to be lost by dividing everything by 2 still applies when there are lots of measures, so I shall assume that a sensible government has taken that point on board to the point where the direct damage from the disease is insignificant compared with the damage caused by the measures. (Just to be clear, I certainly don't mean that lives lost are insignificant, but I mean that the number of lives lost to the disease is significantly smaller than the number lost as an indirect result of the measures taken to control its spread.) Given this assumption, I am free to concentrate just on the damage due to the measures M_i, so this is what I will try to minimize.

The total damage across a full cycle is $\sum_i \eta_i t_i$, so the average damage, which is what matters here, is

$$\frac{\eta_1 t_1 + \eta_2 t_2 + \cdots + \eta_r t_r}{t_1 + t_2 + \cdots + t_r}.$$

We don't have complete freedom to choose, or else we'd obviously just choose the smallest η_i and go with that. The constraint is that the growth rate of the virus has to end up where it began: this is the constraint that $\lambda_1 t_1 + \cdots + \lambda_r t_r = 0$, which we saw earlier.

Suppose we can find $s_1, \dots, s_r$ such that $\sum_i s_i = \sum_i \lambda_i s_i = 0$, but $\sum_i \eta_i s_i \neq 0$. Then in particular we can find such $s_1, \dots, s_r$ with $\sum_i \eta_i s_i < 0$. If all the t_i are strictly positive, then we can also choose them in such a way that all the $t_i + s_i$ are still strictly positive. So if we replace each t_i by $t_i - s_i$, then the numerator of the fraction decreases, the denominator stays the same, and the constraint is still satisfied. It follows that we had not optimized.

Therefore, if the choice of $t_1, \dots, t_r$ is optimal and all the t_i are non-zero (and therefore strictly positive — we can't run some measures for a negative amount of time) it is not possible to find $s_1, \dots, s_r$ such that $\sum_i s_i = \sum_i \lambda_i s_i = 0$, but $\sum_i \eta_i s_i \neq 0$. This is equivalent to the statement that the vector $(\eta_1, \dots, \eta_r)$ is a linear combination of the vectors $(1, 1, \dots, 1)$ and $(\lambda_1, \dots, \lambda_r)$. In other words, we can find σ, ρ such that $\eta_i = \sigma - \rho\lambda_i$ for each i. I wrote it like that because the smaller λ_i is, the larger the damage one expects the measures to cause. Thus, the points form a descending sequence. (We can assume this, since if one measure causes both more damage and a higher growth rate than another, then there can be no reason to choose it.) Thus, ρ will be positive, and since at least some λ_i are positive, and no measures will cause a *negative* amount of damage, σ is positive as well.

The converse of this statement is true as well. If $\eta_i = \sigma - \rho\lambda_i$ for every i, then $\sum_i \eta_i t_i = \sigma \sum_i t_i - \rho \sum_i \lambda_i t_i = \sigma \sum_i t_i$, from which it follows that the average damage across the cycle is σ, regardless of which measures are taken for which lengths of time.

This already shows that there is nothing to be gained from having more than one measure for the relaxation phase and one for the lockdown phase. There remains the question of how to choose the best pair of measures.

To answer it, we can plot the points (λ_i, η_i). The relaxation points will appear to the right of the y-axis and the suppression points will appear to the left. If we choose one point from each side, then they lie in some line $y = \sigma - \rho x$, of which σ is the intercept. Since σ is the average damage, which we are trying to minimize, we see that our aim is to find a line segment joining a point on the left-hand side to a point on the right-hand side, and we want it to cross the y-axis as low as possible.

It is not hard to check that the intercept of the line joining $(-\mu, \xi)$ to (λ, η) is at $\frac{\mu\xi+\lambda\eta}{\mu+\lambda}$. So if we rename the points to the left of the y-axis $(-\mu_i, \xi_i)$ and the points to the right (λ_j, η_j), then we want to minimize $\frac{\mu_i\xi_i+\lambda_j\eta_j}{\mu_i+\lambda_j}$ over all i, j.

Can we describe the best choice in a less formal way?

It isn't completely easy to convert this criterion into a rule of thumb for how best to choose two measures, one for the relaxation phase and one for the suppression phase, but we can draw a couple of conclusions from it.

For example, suppose that for the suppression measures there is a choice between two measures, one of which works twice as quickly as the other but causes twice as much damage per unit time. Then the corresponding two points lie on a line with negative gradient that goes through the origin, and therefore lies below all points in the positive quadrant. From this it follows that the slower but less damaging measure is better. Another way of seeing that is that with the more severe measure the total damage during the lockdown phase stays the same, as does the total damage during the relaxation phase, but the length of the cycle is decreased, so the *average* damage is increased.

Note that I am not saying that one should always go for less severe measures — I made the strong assumption there that the two points lay in a line through the origin. If we can choose a measure that causes damage at double the rate but acts three times as quickly as another measure, then it may turn out to be better than the less damaging but slower measure.

However, it seems plausible that the set of points will exhibit a certain amount of convexity. That is because if you want to reduce the growth rate of infections, then at first there will be some low-hanging fruit — for example, it is not costly at all to run a public-information campaign to persuade people to wash their hands more frequently, and that can make quite a big difference — but the more you continue, the more difficult making a significant difference becomes, and you have to wheel out much more damaging measures such as school closures.

If the points were to lie on a convex curve (and I'm definitely not claiming this, but just saying that something like it could perhaps be true), then the best pair of points would be the ones that are nearest to the y-axis on either side. This would say that the best strategy is to alternate between a set of measures that allows the disease to grow rather slowly and a set of measures that causes it to decay slowly again.

This last conclusion points up another defect in the model, which is the assumption that a given set of measures causes damage at a constant rate. For some measures, this is not very realistic: for example, even in normal times schools alternate between periods of being closed and periods of being open (though not necessarily to a coronavirus-dictated timetable of course), so one might expect the damage from schools being 100% closed to be more than twice the damage from schools being closed half the time. More generally, it might well be better to rotate between two or three measures that all cause roughly the same rate of damage, but in different ways, so as to spread out the damage and try to avoid reaching the point where the rate of one kind of damage goes up.

Summary of conclusions.

Again I want to stress that these conclusions are all quite tentative, and should certainly not be taken as a guide to policy without more thought and more sophisticated modelling. However, they do at least *suggest* that certain policies ought not to be ruled out without a good reason.

If adaptive triggering is going to be applied, then the following are the policies that the above analysis suggests. First, here is a quick reminder that I use the word "measure" as shorthand for "set of measures". So for example "Encourage social distancing and close all schools, pubs, restaurants, theatres, and cinemas" would be a possible measure.

1. There is nothing to lose and plenty to gain by making the triggers (that is, the infection rates that cause one to switch from relaxation to suppression and back again) low. This has the consequence that the triggers should be set in such a way that the damage from the measures is significantly higher than the damage caused by the disease. This sounds paradoxical, but the alternative is to make the disease worse without making the cure any less bad, and there is no point in doing that.
2. Within reason, the cycles should be kept short.
3. There is no point in having more than one measure for the relaxation phase and one for the suppression phase.

4. If you must have more than one measure for each phase, then during the relaxation phase the measures should get more relaxed each time they change, and during the suppression phase they should get less strict each time they change.
5. Given enough information about their consequences, the optimal measures can be determined quite easily, but doing the calculation in practice, especially in the presence of significant uncertainties, could be quite delicate.

Point number 1 above seems to me to be quite a strong argument in favour of the hammer-and-dance approach. That is because the conclusion, which looks to me quite robust to changes in the model, is that the triggers should be set very low. But if they are set very low, then it is highly unlikely that the enormous damage caused by school closures, lockdowns etc. is the best approach for dealing with the cases that arise, since widespread testing and quarantining of people who test positive, contacts of those people, people who arrive from certain other countries, and so on, will probably be far less damaging, even if they are costly to do well. So I regard point number 1 as a sort of reductio ad absurdum of the adaptive-triggering approach.

Point number 2 seems quite robust as well, but I think the model breaks down on small timescales (for reasons I haven't properly understood), so one shouldn't conclude from it that the cycles should be short on a timescale of days. That is what is meant by "within reason". But they should be as short as possible provided that they are long enough for the dominant behaviour of the infection rate to be exponential growth and decay. (That does not imply that they should not be shorter than this — just that one cannot reach that conclusion without a more sophisticated model. But it seems highly likely that there is a minimum "reasonable" length for a cycle: this is something I'd be very interested to understand better.)

Point number 3 was a clear consequence of the simple model (though it depended on taking 1 seriously enough that the damage from the disease could be ignored), but may well not be a sensible conclusion in reality, since the assumption that each measure causes damage at a rate that does not change over time is highly questionable, and dropping that assumption could make quite a big difference. Nevertheless, it is interesting to see what the consequences of that assumption are.

Point number 4 seems to be another fairly robust conclusion. However, in the light of 3 one might hope that it would not need to be applied, except perhaps as part of a policy of "rotating" between various measures to spread the damage about more evenly.

It seems at least possible that the optimal adaptive-triggering policy, if one had a number of choices of measures, would be to choose one set that causes the infections to grow slowly and another that causes them to shrink slowly — in other words to fine tune the measures so as to keep the infection rate roughly constant (and small). Such fine tuning would be very dangerous to attempt now, given how much uncertainty we are facing, but could become more realistic after a few cycles, when we would start to have more information about the effects of various measures.

One final point is that throughout this discussion I have been assuming that the triggers would be based on the current rate of infection. In practice of course, this is hard to measure, which is presumably why the Imperial College paper used demand for intensive care beds instead. However, with enough data about the effects of various measures on the rate of spread of the virus, one would be less reliant on direct measurements, and could instead make inferences about the likely rate of infection given data collected over the previous few weeks. This seems better than using demand for ICU beds as a trigger, since that demand reflects the infection rate from some time earlier.

This entry was posted on March 28, 2020 at 5:48 pm and is filed under News. You can follow any responses to this entry through the RSS 2.0 feed. You can leave a response, or trackback from your own site.

49 Responses to “How long should a lockdown-relaxation cycle last?”

Richard Baron Says:
March 28, 2020 at 6:26 pm | Reply
I wonder how growing immunity plays into this. The answer may be obvious to some, but it isn’t to me. Two thoughts:

If we eventually get immunity in a substantial proportion of the population (that is, if you do not succeed in keeping the infection rate small all the way through), then the rates of infection in all phases should drop because before Fred is immune, the virus could pass from Joe through Fred to Brenda, but after Fred has become immune, that path will be closed.

On the other hand that fall in rate of infection may not matter, so that all your conclusions still hold, if we are only trying to decide what to do in the forthcoming cycle (rather than trying to decide at the beginning what the plan will be for all cycles). It will just be that the decision for the forthcoming cycle might be a bit different from the decisions for preceding cycles.

Sylvestra Says:
March 28, 2020 at 6:27 pm | Reply
I follow the information on covid since January and the thing that surprised me most was how little reliable data we have during the pandemic. So the most problematic assumption is that we can determine all these values well enough.

Andreas Thom Says:
March 28, 2020 at 6:31 pm | Reply
I guess what changes the dynamics is that the growth should be according to a logistic model that takes into account that the rate of infections is also proportional to “population size minus number of infected people”. I would guess that this means that in finitely many cycles one would reaches a final state where (after infinite time) finally everybody will become infected.

Richard Baron Says:
March 28, 2020 at 6:49 pm
That sort of consideration seems to have been taken into account in the Oxford paper (still provisional and over-hyped by journalists) available here: