QUESTIONS ABOUT CORONAVIRUS POLICY

TIMOTHY GOWERS

Disclaimer. I am not an expert on epidemiology, so what I am presenting here are a few thoughts and back-of-envelope calculations based on very simple models that could be oversimplified in important ways. I have tried to indicate areas of uncertainty (by which I mean my own personal uncertainty – I would hope that advice is being sought from people who know a lot more about those areas) wherever possible, but if a simplified model leads sufficiently clearly to a given conclusion, then for the conclusion to be invalid, the model has to be completely wrong, rather than merely a bit crude in some of its details.

I would also like to make clear that nothing I write here will be new to Chris Whitty or Sir Patrick Vallance, or other scientific advisers to the government. However, the conclusion I am drawn to is different from what I understand current policy to be, so I hope that presenting my thoughts will still be useful, particularly given how rapidly evolving the situation is. On general grounds it seems better to expose current policy to constant scrutiny and challenge than to fix that policy once and for all.

1. The herd-immunity policy

My understanding of this policy is that if enough people get coronavirus, then the percentage of susceptible people drops to the point where the "r value" (the average number of people each infected person infects) drops below 1, after which the disease dies out exponentially quickly. If the current value of r is 2.5 (I don't know the correct value, and of course it will change in response to social-distancing measures, so this is illustrative only) then one would expect the fraction of susceptible people to need to drop to around $1/2.5 = 40\%$, and therefore that 60% of the population would have to get the disease. I have seen this figure mentioned by Sir Patrick Vallance, and I would guess that a more sophisticated version of the above calculation lies behind it.

2. Possible criticisms

I see three potential criticisms of this policy, which I shall present in strongly increasing order of seriousness.

2.1. **There is evidence that some people can be reinfected.** I know very little about this, but have read reports that some people appear to have recovered and then become ill again. This would appear to be problematic for the herd-immunity policy, but if it applies to only a small fraction of people who get coronavirus, then it just means that a slightly higher percentage of the population needs to get the disease to achieve herd immunity. So this, though a little worrying, does not seem to invalidate the policy.

2.2. **We do not know that people who get coronavirus have lasting immunity.** It might be that getting coronavirus is like getting a cold or flu: perhaps it protects you against the current strain, but leaves you vulnerable when a new strain comes round in some future year.

This seems a more significant potential downside to the herd-immunity policy. I won't say much about it because (i) I don't know enough about the medicine and (ii) I'm sure this point is not lost on the government's advisers. One possible argument is that if the disease does become a flu-like disease that comes back in modified form each year, it is likely that in future years (maybe not next year, but probably by 2022) vaccines will have been developed, and herd immunity can be achieved that way. And even if getting coronavirus leads to only partial protection against future strains, that would still lead to a more resilient population and a slower-growing disease.

2.3. **We do not have the resources to implement the policy in a non-disastrous way.** The most serious problem with the herd-immunity policy, it seems to me, is that we can't infect 60% of the population in a matter of months without overwhelming the hospitals and having to let a very large number of people die untreated.

Here is the very basic way I arrive at that conclusion. The population of the UK is 66 million, and 60% of that is 40 million.

What percentage of those would need intensive care? I've seen figures such as 10%, but let's assume that we manage to bring this down by, for example, isolating people above a certain age, so that on average the people who get it are less vulnerable. And let's also allow for the fact that the true figure may be lower if we're lucky and a lot of people get the disease without our ever knowing about it. So I'll go for a very optimistic figure of 2% instead. That means that 800,000 would need intensive care for the policy to work.

How many of those would actually get intensive care? The figures I see are that we have around 4000 ICU beds at the moment. Most of these are of course in use already, but if we ignore that, we find that each bed must be used for 200 people on average. So even if

each person needed a bed for a week only, it would take four years to get through them all.

I got to that four-year figure with a series of extremely optimistic assumptions, so I think the true figure would be a lot higher.

The situation would change if we could organize a rapid expansion of the number of ICU beds. Here I am once again in an area I know very little about, but from what I have read, this is very unlikely to be possible in anything like the numbers that would be required: we would need a huge number of extra ventilators, and a huge number of extra doctors, nurses and hospital staff, all taking significant risks with their own health. We might be able to improve provision significantly, but surely not by anything like the factor of ten or more that would be required.

I conclude from this calculation that if the aim of the current policy is to achieve herd immunity by next winter (which is what Sir Patrick Vallance seems to be saying) then the cost of that policy is that only a small fraction of people who get the disease seriously will receive treatment, let alone adequate treatment. So it's not all that different, in percentage terms, from doing nothing.

3. What is the alternative?

If we decided that the cost of implementing the herd-immunity strategy was too high, the other obvious strategy is to try to eliminate the disease by adopting extreme social-distancing measures, as countries such as Hong Kong and Singapore have done.

The obvious disadvantage to this is that such measures are economically very damaging. Here is another point where I am out of my depth: I don't know how to estimate the economic damage, and in particular to assess how many lives it would cost.

A second obvious disadvantage is that if this policy is adopted, then there is no obvious end in sight: the moment the social distancing stops, the disease resumes its exponential growth.

Here is one way this policy might work in practice. If sufficiently strong social-distancing measures are adopted for the r ratio to be less than 1, then the disease will (after a time lag of a week or two) shrink exponentially quickly, so after a few weeks it will be down to a very small number of cases. At that point one could relax the social-distancing measures – in particular the more damaging ones – which would probably cause the number of cases to edge up again, or perhaps even grow quite rapidly. Then at a certain point the social distancing could be reintroduced. Perhaps we would end up in a situation where we had to

take a couple of months off every six months or so (or one month in three, or something of the kind). The hope would be that a vaccine would be found before we had to go through too many of these cycles.

My instinct tells me that a policy of this kind would be less costly than the herd-immunity policy, but here again there are important caveats. The first is that, as I have said, I have made no attempt to understand the costs of shutting down large parts of the economy for weeks or months (though I would make the point that this is likely to happen even if the herd-immunity policy is applied). The second is that how one assesses the cost depends on the value one attaches to human life, and different people will have different attitudes to that, though surely we all agree that it is very high. (If I just think of my own selfish interests, I'm 56, so if I have a 60% chance of getting the disease and a 1.3% chance of dying from it if I get it, then I've got a one in a hundred chance of being killed by the current policy, and the same applies to many of the people I hold most dear – I'd vote for a lot of economic hardship to avoid that.)

4. If one adopts a Hong-Kong-style policy, when should it be applied?

This seems obvious to me. If we abandon the herd-immunity policy, then drastic social-distancing measures should be adopted immediately. Let me try to justify this.

Suppose that the number of cases doubles every four days. Then if we wait four days before applying the policy, the number of cases doubles, so twice as many people die. If extreme social-distancing measures lead (after a lag) to the number of cases halving every n days, then we have to add n days to the length of time it will take for the number of cases to become sufficiently low for it to be appropriate to relax those measures. So there is no downside to starting sooner – the total number of people who get the disease will be a lot smaller and the length of time the measures have to be applied will be reduced.

This would need to be explained to the general public, so that they understood why the measures were not disproportionate. A key point to get across here is the lag: for example, if the disease doubles every four days and it typically takes five days for symptoms to appear, then there are more people walking around with the disease who are not yet symptomatic than there are people who have symptoms.

5. Conclusion

As far as I can tell, there are only two reasonable approaches: a brutal herd-immunity approach and the extreme social-distancing approach (the latter needing to be applied

repeatedly). In particular, there is no "middle approach" where we get to herd immunity in a controlled way – to achieve the control we would have to apply a large amount of social distancing anyway, and for much longer than would be needed if we went for a pure social-distancing approach.

Since the two approaches lead to very different actions right now, we are more or less forced to choose which we are going to do, and to choose as soon as possible. The herd-immunity approach leads to a huge number of unpleasant deaths, so there would have to be a very good reason to implement it. One potential reason is that it is going to happen anyway, but the experience of some countries indicates that the extreme social-distancing approach can work in a reasonably short time, so that reason does not convince me.

I am therefore led to think that repeated phases of extreme social distancing, while still very damaging, will be much less damaging than letting hundreds of thousands die, which will almost certainly involve a great deal of social distancing and economic damage anyway.