This post works through an exercise in how to use and interpret models of disease spread. Here are the takeaways for policy analysis:

- 1. You need to know about a model's sensitivity. Particularly in settings where the specific numbers really matter, such as forecasting how many hospital beds will be needed, it's important to take into account the range of possibilities.
- 2. You need to understand the assumptions in the model. What is the model assuming about the impact of warm weather? What time period does it cover? What choices have been made about disease parameters?
- 3. Your strategies need to take uncertainty into account. This means that a lot may depend on how precautionary you want to be. It also means that you have to be able to adjust your strategies as we learn more about the disease.

Climate policy depends on models to forecast the benefits of interventions. Similarly, policy decisions regarding the COVID-19 epidemic also form the basis of planning. We actually understand the climate system much better than we understand the pandemic at this point. This makes it even more important to have a deeper understanding of what models mean.

Running the Exercise.

In climate policy, one of the key parameters is called climate sensitivity, which indicates how much global temperature change results from adding a given amount of carbon to the atmosphere. There are a number of important parameters in a model of the coronavirus, but one of the most important is called R0 (also called the basic reproduction number) which tells how contagious a virus is. R0 is the number of new victims who will catch the disease directly from a single infected individual, assuming there's no natural immunity. We only know the approximate R0 for the coronavirus. An important question in interpreting models is how much the model's projection's are affected by changes in the figure for R0 used in the model.

For this exercise, I used an interactive model on the *NY Times* website that illustrates the issue clearly. I used the model to construct a table showing what happens to the number of deaths with various precautions. I ran the numbers for three different settings of R0, all of them in the ballpark of 2.5 (about what many experts think is the right number). I checked for results with two different levels of policy intervention, one relatively modest and one quite stringent. The "modest" level of intervention that I chose includes two weeks of what the model calls moderate social distancing requirements. The "stringent intervention" means four weeks of strong social distancing plus a lot of testing. What's important in the

table is not the specific predictions about the number of deaths, but how they vary between columns.

I repeated this exercise several times in order to try to be sure that I wasn't screwing up the numbers. For some reason, I got slightly different figures each time, but they were all close to these results.

There are also other parameters about the disease that can be varied: how sensitive the virus is to warmer weather, what percentage of people require hospitalization, and what percentage will die even with treatment. I didn't try to assess the model's sensitivity to those other parameters but I invite you to try that for yourself.

Intervention Strength	R0=2.3	R0=2.5	R0=2.8
Modest	676,000 deaths.	1.4 million deaths	1.8 million.
Stringent	43,500 deaths.	247,000 deaths.	1.1 million

Two things are striking about this table. First, the total number of deaths is extremely sensitive to R0. And second, so is the relative impact of interventions. With an R0 of 2.5, the stringent intervention causes the number of deaths to fall more than eighty percent. With an R0 of 2.8, the decrease is only about forty percent. In absolute terms, with the lower R0, the intervention saves well over a million lives, while with the higher one, it's only 700,000.

Interpreting the Results.

How should we think about this model's sensitivity to the choice of R0? At least part of it probably reflects reality and provides some real insight into the underlying dynamics governing public health interventions. Because coronavirus spread is an exponential process, differences get amplified. An R0 of 2.3 means that one infection today will result in twelve infections after we go through three new "generations" of infection. (That's probably over about two weeks). But an R0 of 2.8 means twenty-one new infections, almost twice as many. Moreover, the mortality rate is sensitive to surges, because hospitals get overwhelmed and cannot treat everyone. A higher R0 means that the surge is bigger and

happens faster. A higher R0 also means that temporary measures that slow the disease have shorter effective periods because it springs back faster.

There's another feature of the model that could be amplifying the effect of R0. The model runs only from January to late October, and it assumes that the disease slows down during warm weather. Thus, the periods of rapid spread (Spring and mid-Fall) are limited. Given this limited time period, it matters a lot how quickly the disease spreads. To check on this effect, I tried turning up the parameter for "impact of warm summers" from moderate to high. As I expected, this made the results much more sensitive to the choice of R0.

Even taking all this into account, this seems to be a very high level of sensitivity that may reflect what I'm sure is a very stripped-down model. It can't be easy to design a model that thousands of people can run on-line to get quick results. I also looked at the highly respected Imperial College model quickly for comparison. While it did show significant sensitivity to R0, it didn't seem to be nearly as dramatic.

Conclusion

Even if the results from running the *Times*'s table-top model are exaggerated, there are some important lessons. First, the model does help us understand the dynamics of the process better, including such things as the importance of any summer slowdown. Second, while the exact degree of sensitivity to R0 may be lower in more elaborate models, it still seems to have a significant impact. And third, it's not enough to know a single projection from a model, you also have to understand the degree of uncertainty.

In a way, I hate talking about model sensitivity because some people will seize on any uncertainty as a justification for inaction. That's a stupid if not immoral response.

As in climate policy, uncertainty is not our friend. Yes, it's possible that the problem will turn out to be less serious than we now expect. But it's also possible that it will turn out to be much worse.