Perfect Mixing in COVID Models

Bob Eberlein isee systems

Abstract

Most models in System Dynamics, and many in epidemiology more broadly, use aggregate difference of differential equations which means that all stock values will have perfect mixing. While it is true that this is not what happens in the spread of a disease, the important question is whether this matters for the use of models in understanding and the methods used to compare model results with measured data. In this paper perfect mixing models are contrasted with conveyor-based models in which material enters and leaves without mixing. These two types of models are compared against the outcome of an agent-based simulation which is taken as the ground truth. In general, the conveyor base model is more like the ground truth model both in terms of local behavior and the interpretation of the underlying parameters. However, the two model types are closer to one another than they are to the agent-based models. This suggests that straight comparison of measured data to model results for aggregate models should be approached with a great deal of caution.

Introduction

We are all aware of both the dangers and advantages of using perfect mixing in representing systems that, by their very nature, are not mixed. This is a topic that has been revisited in various forms in both System Dynamics and other quantitative fields. Each time it is addressed, the answer is invariably that whether it matters depends on the purpose of the modeling work.

When COVID was starting to be recognized as a pandemic, and the initial data on infections and mortality became available, I was struck by the disease duration, as well as the surprising variability in symptom severity. Because perfect mixing pushes the large majority of change in state to the early part of any delay I was concerned that this would cause significant distortions in both representing the epidemic spread and, more importantly, the impact of interventions on that spread.

In this paper, I go through three models that represent the same problem – an agent based stochastic model, a perfect mixing deterministic model, and FIFO conveyor based deterministic model. I show that the parameters that generate similar behavior in the agent-based and conveyor based models generate significantly distorted behavior in the perfect mixing model. Most of this difference can be compensated for by adjustment of the underlying epidemiological parameters. I then try to address the question of whether there is a downside to making such adjustments beyond the simple interpretation of what they mean. This is done by calibrating the conveyor and perfect-mixing models to output of the agent-based model.

Ground Truth

For this paper we are taking the ground truth as an agent-based model of the spread of infectious disease, where the parameters are clearly interpretable as the on-average values for individual realizations of infection and disease states. That model itself is simple, and done in Stella:



The main chain is the same as that of any standard SIR model, except the entries in the chain are individuals and the stocks thus represent states for an individual. An individual can only be in one state at any time, and thus this formulation is effectively a Markov Chain based on the probability of moving from one state to another.

The standard positive feedback loop is clearly visible in this model, though the balancing loop is not (the balancing loops is actually the same loop because the infected state turning off lowers infected contacts). The model equations are less elegant than they would be for a standard System Dynamics models to accommodate both the necessary stochastic nature of this type of model, and the manipulation of the array elements representing individuals.

Individual contacts, arrayed by contact and individual are determined by:

```
IF contact <= effective_contact_rate*DT THEN INT(UNIFORM(1,
SIZE(Individual)+0.9999999) ELSE 0
```

That is, every individual randomly contacts a random selection of every other individual in the population based on the underlying contact rate (all individuals have the same contact rate).

To determine whether the contact is with an infected individual we have as the equation for infected contacts (again arrayed by individual and contact)

```
IF individual_contacts > 0 AND
Infected_individual[individual_contacts] > 0 THEN 1 ELSE 0
```

Which simply says that if the individual being contacted is infected it is 1, otherwise 0.

Given the number of contacts with infected individuals the equation for infections from contacts is given by

```
IF infected_contacts > 0 AND UNIFORM(0, 1) < .infectivity THEN 1 ELSE 0 \,
```

The variable recovery time is set at the time of infection by adding 14 days to the current time.

The other equations should be clear from inspection, and the full model is included in the supplementary materials. The stochastic functions in that model also have seed values to allow replication of results.

This model is very simple and could include more variety. It was constructed in its current form as a comparator to address this question of how much difference the perfect mixing assumption makes when modeling COVID.

Model Results

Starting with a single infected individual and the parameters noted on the model diagram, the following indicates the range of outcomes one might expect to see. This is with a population of 10,000 people.





There is, of course, a good deal of variety on timing, and three of the runs actually did not result in any disease spread (R0 in this model as configured is 2.52). This is a clear distinguishing characteristic of any individually based model (necessarily stochastic) relative to an aggregate model (often deterministic), and consistent with the emphasis on early containment for emerging infectious diseases. In the real world when stopped early, diseases can effectively become extinct, as was the case for SARS and MERS, the two most closely related diseases to COVID.

We will be using Run 8 from the above outcome as a basis for comparison with our different mixing models. It is relatively early in timing, and at 25 days has a 7 day average of new infections of about 5 (per 10,000) which makes it useful as a basis for comparison with the other models at that time. This is not a big number, but it is consistent with the early stage of the pandemic in the US and other countries when everything was growing, and no turnaround was in sight (peak infection rates of ½ to 1 person per thousand were common during the first wave of COVID).

Perfect Mixing

The perfect mixing model is the standard SIR model:



In this case the positive and negative loops are clearly visible.

No Mixing

The model without mixing uses a conveyor instead of a normal stock for the number of people infected:



There were elaborations of these models developed that use a distribution of infection duration available from the author on request. The added complexity does not change the basis outcomes described here (except the extinction result) so the simpler form was used.

Behavior Comparison

Recognizing that there are a range of outcomes from the agent-based model, it is still possible to compare the basic patterns of behavior for the different model formulations. Assuming the same parameters for contact rate and the initial number of infected people (1) we can compart results of the three model configurations.



Qualitatively these are all similar, but there are big differences in timing and, as a consequence, the peak number of people becoming infected.

Looking at only the first 25 days:



We see the type of behavior that prompted this entire line of investigation. The perfect mixing model is starting off very slowly, whereas the version using conveyors seems like a better approximation of the agent-based result. This relatively fast pickup is because there are more people in the infected state when they stay resident for the full 14 days instead of draining out exponentially:



Calibrating to the First 25 Days

Of course, the fact that the same contact rate number seems to get us closer to the expected behavior when using conveyors does not tell us much about what we can learn from using aggregate models. This is because we almost always have to make adjustments to model parameters so that they have a reasonable interpretation when used in any given model.

A typical approach to calibration I have used is to adjust the initial number of infected people (something that is not truly observable), and the underlying contact rate (something that is not observable) so that the model and data largely line up. Taking the agent-based model as the data against which to replicate this yields:

	Not Mixing.init infected	Perfect Mixing.init infected	contact rate not	contact rate perfect	Payoff
Starting at	1	1	40	40	

After					
522	1.39155495889	0.962008961707	39.4786277422	50.0942166707	2.84563230278
runs					

The contact rate for the conveyor based model decreased slightly while that for the perfect mixing model increased a great deal. Conversely, the initial number of infections went up more for the conveyor based model. Nonetheless, relative results look quite similar and are reasonable (we do not expect to be able to capture the jump from 0 observed in the agent based model):



In fact, the two aggregate models are almost indistinguishable. This remains true running the model out to day 100:



There is a small difference between the two aggregate representations, but both are dramatically different from the agent-based model. Of course, all of these differences dwarf the numbers before time 25 because of the speed at which they are increasing.

The main question I was trying to pose when doing this work, was what difference it would make to social distancing for a limited amount of time. To test this, we use the parameters from the calibration, and turn on the intervention at time 25, turning it off 14 days later. The results shown to time 50 are similar for the two aggregate models.



However, if the effectiveness of the intervention is higher, there is substantial divergence between the two approaches:



And setting it to 100% effective for 14 days (the dream of every lockdown):



Though striking, this outcome is somewhat contrived as the diseased duration is fixed in both the agentbased and conveyor-based model. The perfect mixing model, on the other hand, has a negative exponential distribution of disease duration. Changing the disease duration to a distribution (stochastic on the agent-based model and deterministic for the conveyor-based model) does not substantially change any but the last outcome.

Conclusion and Possible Extensions

All models represent simplifications of reality and it is the demonstration of fitness to purpose that is paramount in deciding whether of not to make use of them. Still, one of the biggest dangers in doing prospective modeling work is having something that can do well at replicating observed behavior while failing miserably when used for its intended purpose.

In this paper I have demonstrated that moving away from perfect mixing in representing infectious disease dynamics does improve the fidelity of the model, at least locally. The difference in relative performance of the two types of models is, however, dwarfed by their divergence from the agent-based model which is likely to be a better representation of the underlying dynamics. Using calibration as a means of adjusting parameters when fundamental structural differences exist between the true data generation process needs to be approached with extreme caution.

The models in this paper have all been extremely simple. However, the demonstrated divergence is something that should be investigated for more complete models. It is incumbent on all those doing serious modeling work in this area to ensure that the ways in which models are being compared with data are consistent with the purposes for which the models are being used.