An Acute Hospital Demand Surge Planning Model for the COVID-19 Epidemic using Stock-and-Flow Simulation in Excel: Part 1.

Authors:

Simon Dodds, Health Care Systems Engineer, Birmingham, UK.
✉ simon.dodds@saasoft.co.uk

Declan T. Bradley, Consultant in Public Health Medicine and Clinical Lecturer in Public Health, Belfast, UK.
✉ declan.bradley@hscni.net

Sponsors:

Mike Davidge, Phil Debenham, Martin Dowdall, Ken Fullerton, Chris Jones, Peter Kidd, Ellen Pirie, Delia Pudney, Kate Silvester, Lee Talbot, Lesley Wright.

Abstract

An emerging infectious disease, COVID-19, has spread from an infection cluster in Wuhan City, Hubei Provence, China, to create a global outbreak (pandemic) and, at the time of writing, has infected more than 180,000 people and caused 7,000 deaths.

Epidemiological modelling shows that non-pharmacological interventions (NPI) such as limiting movement of people and social distancing offer the hope of flattening the peak of load on the health care system of each country, and of reducing mortality.

Planning the required acute hospital bed capacity to absorb the expected epidemic surge requires prediction of the magnitude and timing of the local demand on an acute hospital.

This essay describes the science underpinning epidemiological modelling and demonstrates that a basic system dynamic (or stock-and-flow) model of an epidemic can be implemented using widely available spreadsheet software such as Excel.

Using evidence known to date, this basic model illustrates comparable system behaviour to the more sophisticated simulations that are being used to guide the national response to this rapidly developing threat. Specifically, that immediate, nationwide adoption of behaviours such as social distancing could significantly mitigate the effects.

This first model could serve as a useful informational and educational resource at a time of great uncertainty, and further development is now underway to calibrate the model so it can be used to support tactical contingency planning. (223 words)

Keywords

Health care; Infectious disease; Epidemiology; Public health; COVID-19; Pandemic; Contingency planning; System dynamics; Simulation model; Stock-and-flow; Excel; Health care systems engineering (HCSE).

Context

In eleven weeks, the emerging infectious disease COVID-19 has spread from an infection cluster in Wuhan City, Hubei Provence, China, to a global outbreak (pandemic) that, at the time of writing, has infected more than 180,000 people and caused 7,000 deaths.

Public health agencies sought to contain the infection by resource-intensive tracing and isolation of exposed people. Continuing transmission demonstrates that the outbreak has not been contained. The UK now aims to delay the effects of the outbreak, with the goal of spreading the outbreak over a longer time, with a lower peak and reduced overall mortality. In a health and social care service that was already under pressure, an impending sudden increase in demand for health care means that service managers, clinicians, provider organisations and commissioners will need to make rapid decisions about continuing elective work and future bed capacity, working patterns and staff resources. Community social care services will be faced with outbreaks in residential facilities and domiciliary care services being affected by staff absences, which will affect hospital flow.

Control of an outbreak of an infectious agent that spreads from person-to-person requires interruption of the chain of transmission by reducing the number of infectious people (through treatment), reducing the number of susceptible people (through immunisation), and/or reducing opportunities for direct and indirect transmission between infectious and susceptible people. For COVID-19, there is currently no immunisation or specific treatment, so control requires reducing exposure. Risk-reduction may be achieved by numerous means, some of which are specific to the infection's mode of transmission: identification and isolation of affected people and their contacts, environmental decontamination, hygiene, sanitation and social distancing. In the UK during the current epidemic, the response relied initially on identifying and isolating cases and contacts; and motivating personal protective behaviours through public health risk communication. At the time of writing, the United Kingdom (UK) now recommends voluntary 'social distancing' [Ferguson 16-03-2020].

The UK's emergency response is led by the government with advice from the Scientific Advisory Group for Emergencies (SAGE) and other expert groups. The modelling conducted at national level for the 'reasonable worst case' scenario does not aim to guide local understanding of how individual health and social care services, such as hospitals, care homes and community care services, will be affected by the pandemic or how they should respond to the changing situation given their own specific contexts.

The traditional and simplest epidemic model is the *Susceptible-Infectious-Recovered* multi-compartment model, represented in Fig. 1 and despite the structural simplicity, the dynamic behaviour of the SIR model is non-linear, complex and counter-intuitive.

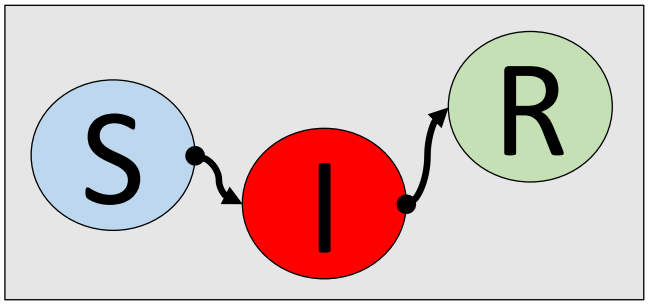

Figure 1. The SIR model for an infectious disease. Individuals move from S (susceptible), sequentially through I (infectious) to R (recovered = non-infectious) as their illness progresses as illustrated by the arrows. The number of people in each compartment influences these flows so the dynamic behaviour of the SIR system can become very complex despite the structural simplicity.

Lessons learned from previous epidemics of influenza, SARS, MERS-CoV, and Ebola, together with computer simulation models are widely used to assist the understanding of the behaviour of infectious diseases and in the pre-emptive planning for a possible future epidemic caused by a novel infectious agent. Local management of an epidemic could also be greatly assisted by simulation models which facilitate planning of novel solutions in a rapidly evolving crisis.

The health and social care system is an example of a complex adaptive system (CAS) and health care systems engineering (HCSE) is an emerging and evolving discipline that draws on the theory, techniques and tools of systems engineering (SE). One such technique is system dynamics (SD) simulation modelling that is already used to assist the diagnosis, design and delivery of improved health care processes.

The COVID-19 epidemic is caused by a novel coronavirus (SARS-CoV-2) for which the global population has an unknown level of background immunity. Though the virus causes a mild, self-limiting respiratory infection in the majority, it can lead to a more severe pneumonia requiring hospital admission for supportive treatment. A small proportion of patients become critically ill and require more intensive support including ventilation, and a significant proportion of severely or critically ill patients do not survive.

Our immune system protects us from pathogens so even when exposed we may avoid becoming infected. But when exposed to a novel infective agent it takes time for our immune system to mount a defence, during which time the pathogen can multiply and spread. This may cause symptoms such as a high temperature (pyrexia), and the interval between becoming infected and exhibiting symptoms is called the *incubation period* (Fig. 2). While we are infected, we may also become infectious (i.e. can transmit the pathogen to others) and the interval between being infected and becoming infectious is called the *latent period*. If the latent period is shorter than the incubation period, then an infectious agent can be unwittingly transmitted before the individual is aware that they are infected or infectious. The emerging evidence is that this novel coronavirus has a latent period that is shorter than the incubation period so can be transmitted unintentionally.

Figure 2. Relationships between the events, states, compartments and named intervals.

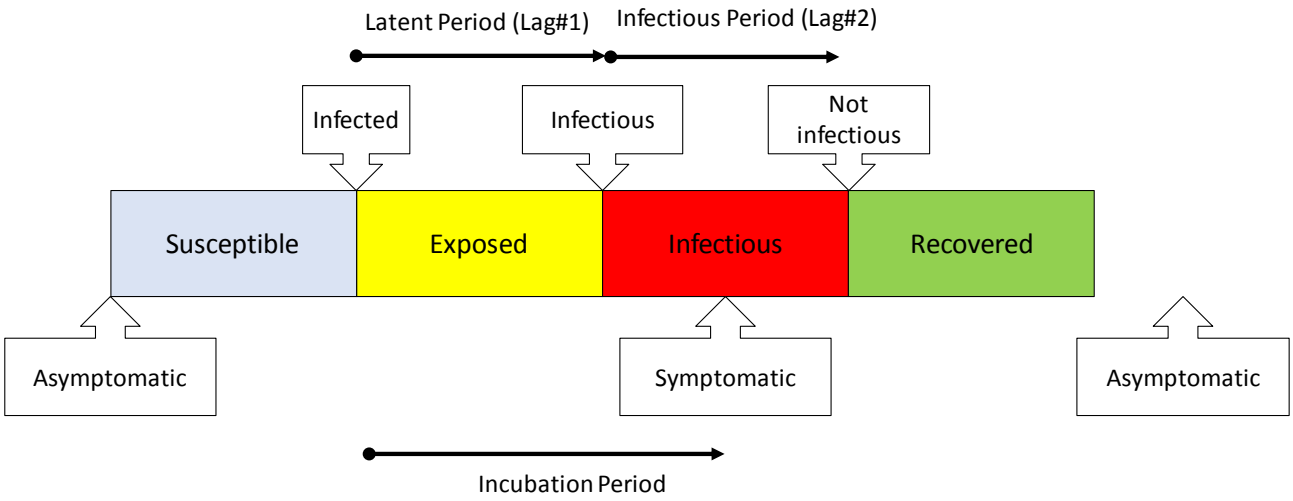

The ability of an infectious agent to spread in a population is also influenced by the number of individuals who have immunity to the pathogen. This population immunity is initially low and increases either as a result of individuals being exposed to the agent, or by prophylactic strategies such as immunisation programmes.

So, the novelty and nature of the COVID-19 virus implies that the combination of factors may create a 'perfect storm' with a potentially rapid growth in the number of severely ill patients that then threatens to overwhelm the health care system causing further collateral problems as patients with other severe illnesses cannot then access services.

Health care systems globally are working to understand the challenge and to prepare for it, and there is an urgent need for predictive tools to assist capacity planning. Specific requirements include planning the number of hospital beds, equipment, staff and consumables required to manage the expected peak of demand; and re-designing the health care processes to make best use of the resources available. Each acute hospital will need to adapt the strategic guidance provided by their government to their local needs and resources, but service managers are usually not trained in the skills required to rapidly design, build, verify and deploy predictive modelling tools. They are, however, generally very familiar with using spreadsheet software such as Microsoft Excel.

Purpose

The purpose of this essay is to inform and educate by visualising the complex dynamics of epidemics using an Excel-based SIR simulator built using systems engineering principles combined with rapid application development (RAD). This is the first component of an adaptable acute hospital capacity planning tool.

Method

A structurally simple, time-dependent, population model for human-to-human infectious disease transmission where recovery confers long-term immunity consists of three compartments: S for the number of susceptible people, I for the number of infectious people, and R for the number recovered (i.e. not infectious) people - hence the name SIR (Fig. 1). In a rapidly spreading epidemic the total population (N) is assumed not to change so at any time **t** after the introduction of the infectious agent:

$$N = S_t + I_t + R_t \qquad \text{Eq. 1}$$

where S_t is the size of the susceptible population at time **t**, I_t is the number of infectious individuals and R_t is the number of patients who have recovered to the point they are no longer infectious or who did not survive so have been removed from the system.

The flow from Susceptible to Infectious at time **t** is given by $\beta * I_t$ where β is the average number of contacts per person times the probability of disease transmission in a contact between a susceptible and an infectious subject. The flow from Infectious to Recovered at time **t** is given by $\gamma * I_t$ and represents infectious patients becoming non-infectious by being isolated, becoming immune or not surviving. If the duration of the infectious period is D, then $\gamma=1/D$ or by re-arranging $D = 1/\gamma$.

In epidemiology, the *basic reproduction number* (R_0) or "R nought" of an infectious disease is the average number of new cases directly generated by one infectious individual in a population where all individuals are susceptible to infection and is given by β/γ. If R_0 is greater than 1.0 the infectious compartment will initially grow exponentially, and this is the necessary condition for an epidemic to develop. The larger R_0, the steeper the exponential growth, and the larger the proportion of the susceptible population that become ill. R_0 is not to be confused with the *effective reproduction number* (eR) which is the number of cases generated in the current state of a population. Emerging evidence shows that R_0 for COVID-19 lies in the range 1.4 to 3.9.

*Figure 3. The set of three differential equations that describe the SIR model [Wikipedia 19-03-2020]. This set of time-dependent equations can be solved numerically for any set of inputs (N, β, γ) and initial values for S, I and R. For an epidemic to start I needs to be greater than zero, and then the size of I will increase exponentially provided there is a large enough proportion of susceptible individuals in the population (S/N) and β>γ (i.e. R_0>1.0). Eventually, the size of the susceptible compartment S falls sufficiently that the rate of recovery (γ*I) starts to exceed the rate of infection (β*I*S/N), and that is when the epidemic peaks. From then on, the size of the infectious compartment I falls progressively to zero at which point the epidemic ends.*

$$\frac{dS}{dt} = -\frac{\beta IS}{N},$$
$$\frac{dI}{dt} = \frac{\beta IS}{N} - \gamma I,$$
$$\frac{dR}{dt} = \gamma I.$$

Exponential growth is difficult to comprehend intuitively, so visualising the time-dependent behaviour of the SIR multi-compartment model is a necessary aid to understanding. A convenient way to do this is to use spreadsheet programme such as Microsoft Excel.

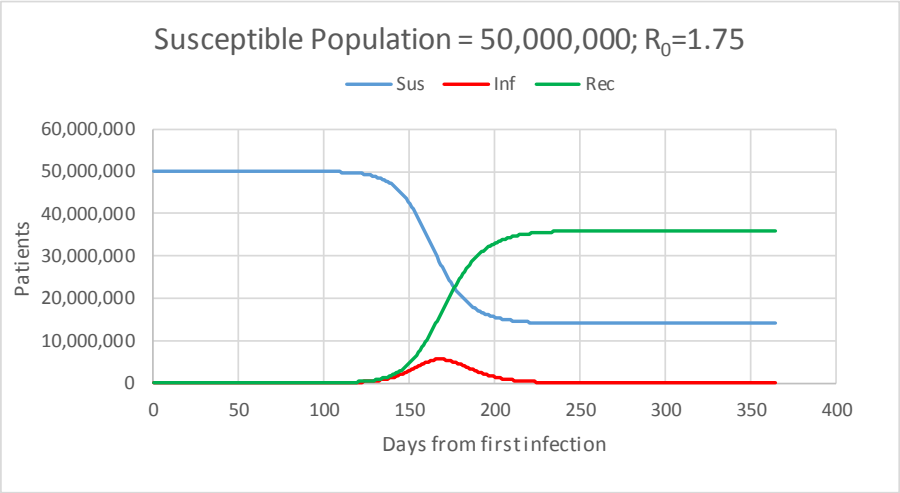

Figure 4. Visualisation of the time course of an epidemic of a novel infectious agent introduced at t=0 into a susceptible population S of 50 million people where R_0=1.75.

The blue line is the size of the susceptible population (Sus), the red line is the number of infectious individuals (Inf), and the green line is the number of patients who are no longer infectious (Rec).

Note the time lag of more than 100 days before the number of infectious cases (Inf) starts to be visible, the accelerating growth of the infectious population, and the high proportion of people that become non-infectious (Rec). The epidemic is self-limiting because eventually the size of the susceptible population (Sus) falls. Note that when the epidemic ends there is still a proportion of the population who are susceptible, in this case about 30%.

One strategy to prevent epidemics is to immunise some of the population so that the ratio S/N falls. The condition for this to prevent an epidemic is β*I*S/N = γ*I and if the proportion of immunised individuals is p then S = N(1-p) and we get β*(1-p)=γ which, substituting with R_0=β/γ can be simplified to:

$$p = 1 - 1/R_0 \qquad \text{Eq. 2}$$

So, for the epidemic illustrated in Fig. 1 where R_0=1.75, the value of p is 1-1/1.75 = 0.43 which implies that only 43% of the population needs to be immunised to prevent an epidemic developing. This is called the *herd immunity threshold* (HIT) and is much lower than the 70% of the population who became ill and then immune when the epidemic ran its natural course.

The value of R_0 is not solely a characteristic of a specific infectious agent, as it is also influenced by the context and environment because the rate of infection is influenced by interventions that reduce the opportunity of cross-infection such as hand washing, wearing personal protective equipment (PPE), or separating infectious and susceptible groups (e.g. quarantine, restriction of movement, social distancing). For some infections it is also affected by other contextual factors such as sanitation and humidity. So, given

that we can intervene to reduce R_0 to contain, delay and mitigate the spread of an infectious agent in a susceptible population, it is important to understand the relationship between R_0 and the size and timing of the peak of illness. We can use the theoretical SIR model to explore this relationship (Fig. 5).

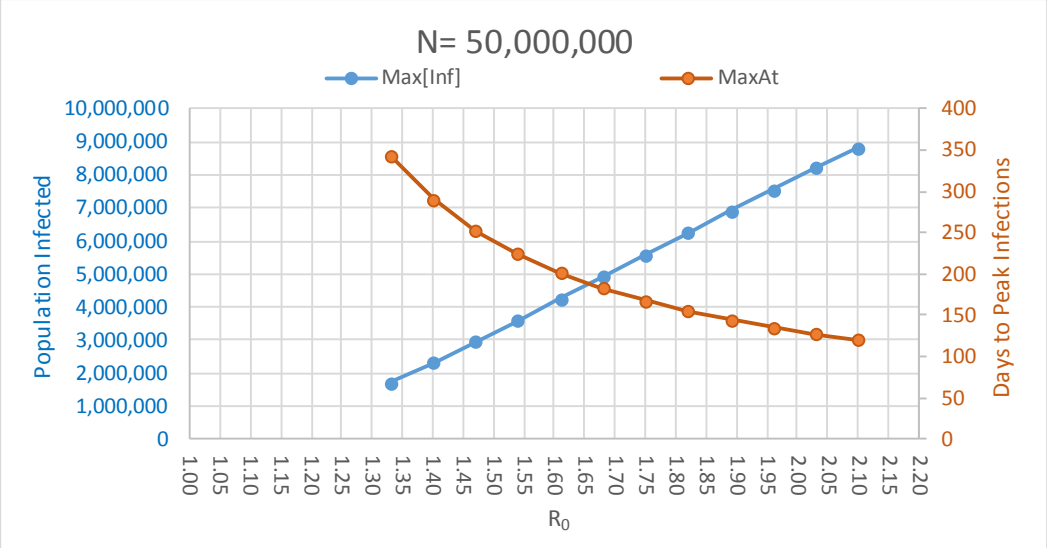

Figure 5. Relationship between the value of R_0 and the number of infectious patients at the peak of the epidemic (blue), the time to peak (brown) for a susceptible population of 50 million.

Note as R_0 increases both the peak and speed of onset of the epidemic increase. So, a rational strategy to lower and delay the peak is to reduce R_0 (i.e. β/γ) and that implies reducing β and increasing γ. Comprehensive contact tracing and isolation is one way to increase γ by removing those individuals from the infectious pool more quickly. Personal infection protection behaviours, such as respiratory hygiene, hand-washing and social distancing, are examples of ways to reduce β.

Experience of the COVID-19 epidemic in other countries indicates that some individuals require hospital treatment, especially those with pre-existing conditions and poor physiological reserve. As health care resources are limited and usually already over-stretched, the additional shock of a large epidemic could tip the whole healthcare system into collapse. So local health care systems will need to be able to act quickly, decisively and effectively and that implies being able to make a reliable prediction of the effect of an action.

Predictive modelling can capture the evolving context and emerging system behaviour and one technique is a system dynamics (SD) or stock-and-flow (S&F) model that can often be implemented using spreadsheet software. The generic sequence used in systems engineering is:

1-Specify 2-Design 3-Build 4-Verify 5-Implement 6-Validate

1. Design Specification

An easy-to-use, Excel-based, interactive simulation tool that can be used to educate health care staff about the counter-intuitive behaviour of epidemics and to assist with estimating the dynamically changing number of acute and critical care hospital beds needed to manage an evolving COVID-19 epidemic.

The design of a stock-and-flow model requires two steps:

2.1 Causal Loop Diagram
2.2 Stock-and-Flow Diagram

2.1 Causal Loop Diagram

A causal loop diagram (CLD) or map is used to illustrate the relationships between the variables in this simplified system and to explain the typical evolution over time of an infectious disease epidemic (Fig. 6).

Figure 6. Causal loop diagram of a typical infective disease. The circles and boxes represent system variables, and the arrows are causal links. A red arrow with a positive sign indicates that the input (cause) moves the output (effect) in the same direction; a green arrow with a negative sign indicates the input moves the output in the opposite direction. A causal loop is any closed path created by following the arrows. A causal loop with an odd number of negative arrows is called a balancing or stabilising loop (B-), otherwise it is a reinforcing or destabilising loop (R+).

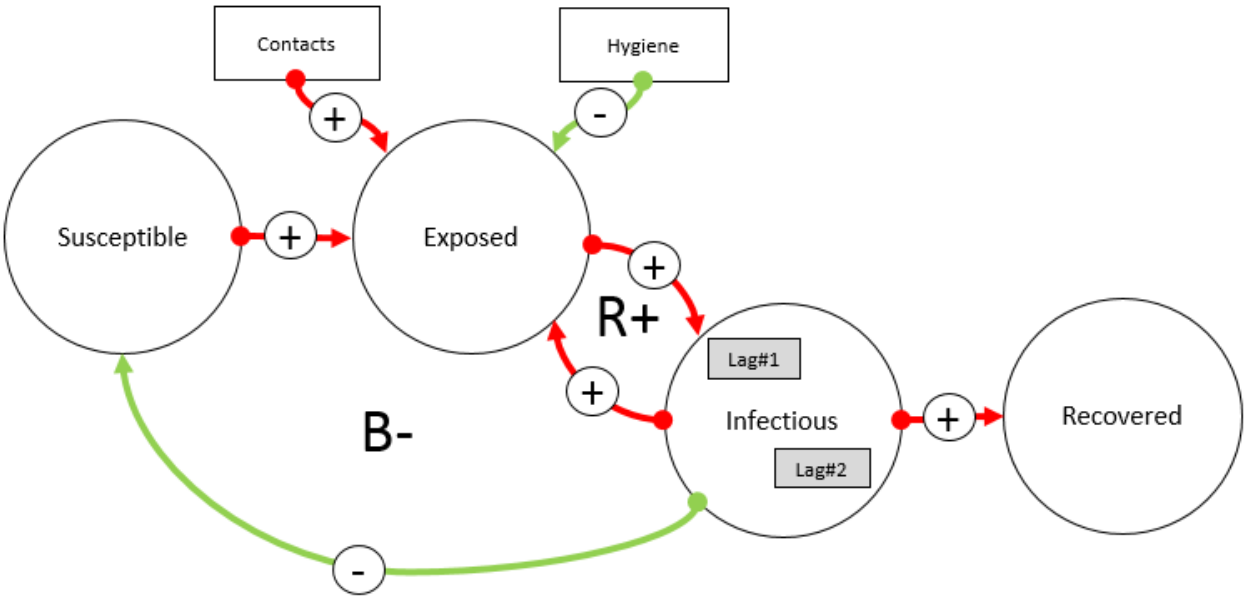

The causal loop diagram for a self-limiting, infectious disease is based on the sequential process of susceptible-exposed-infectious-recovered (SEIR) (Fig. 2). There is a reinforcing loop (R+) between the Exposed and Infectious compartments because infectious patients are the source of new exposures and this causes the epidemic to grow and spread. There is a lag from exposure to becoming infectious that is caused by the delay while the virus replicates and spreads in the infected individual (latent period, lag#1); and a second lag before infectiousness drops as the individual's immune response eliminates the virus (infectious period, lag#2). Infectious patients are assumed to be no longer susceptible which reduces the susceptible population so forms a balancing loop (B-) which will eventually terminate the epidemic. The probability of a susceptible person becoming infectious is determined by the number of person-to-person contacts per day, the transmissibility of the pathogen, and the proportion of infectious people in the population (Figs 2,3). If R_0 is less than 1.0 then the disease will not spread; if the R_0 is more than 1.0 then the number of infected people will increase exponentially (epidemic potential). Eventually, a patient will become non-infectious and move to the final compartment called Recovered. In the case of an epidemic, the pool of susceptible people decreases over time so the opportunity for an infectious person to infect someone else will also fall. This illustrates the dynamic nature of the *effective reproduction number* (eR), and when it goes below 1.0 the epidemic peaks and will then wane as the infectious pool empties.

To make specific predictions of the plausible behaviour in a real system, the causal loop diagram is used to develop a system dynamic or stock-and-flow map that is the blueprint for building a practical simulation tool.

2.2 Stock-and-Flow Diagram

A stock-and-flow diagram or map shows the structure of a system using specific symbols that distinguish between quantities that are stocks and those that are flows; which differ in their units of measurement. A stock is measured at a specific time and represents a quantity existing at that point in time (the S, I and R compartments (Figs 1, 2) are stocks. A flow links two stocks.

Figure 7. Stock-and-flow map of the health care macrosystem used as a blueprint for building an Excel-based simulation model. Stocks are shown as rectangles, flows as black solid arrows with flow control "valves". Note that a flow is one way with the direction shown by the arrow. The control signals are shown as dotted lines with arrows indicating the direction of information flow and the internal flow control logic is indicated by blue-shaded boxes. The prefixed names refer to system variables in the Excel model: w- implies a stock, q- implies a flow, p- implies a probability, n- implies a count, and t- implies a time interval (see Table 1 in the Glossary for descriptions).

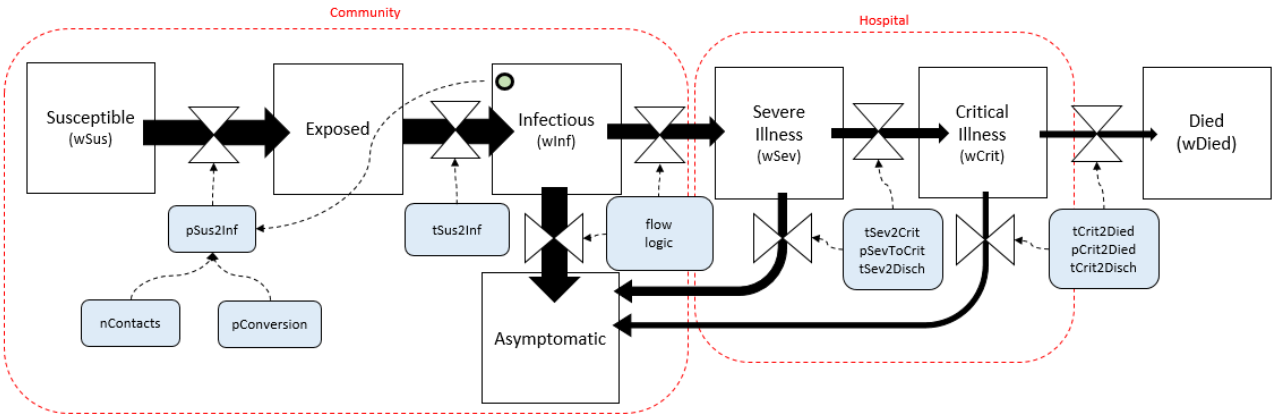

3. Build

It is relatively straightforward to build this system dynamics model in Excel because there is a one-way flow from the starting Susceptible stock (the source) to the final Recovered and Died stocks (the sinks), and because there are no flow loops. Fig. 7 illustrates the alternating sequence of stock-flow-stock and this can be used to structure the Excel model to make it easier to build and verify. Each row in the sheet represents the flow in a time interval (*dt*) and the columns represent the alternating stocks and flows. Stocks are coloured blue and have a w- prefix and flows are coloured orange and have a q- prefix and a name that indicates which stocks the flow links to.

Figure 8. Screen shot of part of an Excel implementation. Time runs from top to bottom and the flow from left to right. The cells in red on the second row indicate the initial values of the stocks. The logic of the model is implemented using formulae in the cells (not shown). There is also an input control area and an output display area. Once the model is verified to be working correctly it is usual to password protect all but the input cells to avoid accidental changes that would invalidate the outputs. See Glossary for definitions.

	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y
1	W/B	wSus	qSus2Inf	wInf	qInf2Rec	wRec	qSev	wSev	qSev2Disch	qSev2Crit	wCrit	qCrit2Disch	qDied	wDied
2	01/01/2020	49999999.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
3	02/01/2020	49999999.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
4	03/01/2020	49999999.00	0.25	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
5	04/01/2020	49999998.75	0.25	1.25	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
6	05/01/2020	49999998.50	0.25	1.50	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
7	06/01/2020	49999998.25	0.31	1.75	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
8	07/01/2020	49999997.94	0.37	2.06	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
9	08/01/2020	49999997.56	0.44	2.44	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

4. Verify

The verification step of the systems engineering process is essential to ensure that unintended errors have not been introducing during the build stage. A series of verification tests are specified as part of the design phase and each represents a scenario with a defined set of inputs and an expected output. The actual output of the model is compared with the expected output and any differences trigger an investigation and correction of the cause. Only when all the verification tests have been passed can the model outputs be trusted. For a model with many inputs there are potentially a very large number of possible combinations, and the design of a comprehensive set of verification tests can be the most difficult part of the process.

Figure 9. Screen shot of part of the verification test table. The required input values are defined for each scenario and typically only one input is changed at a time to facilitate investigation of the cause of a verification test failure. See Table 1 in the Glossary for definitions of system variables.

Acute hospital surge capacity planning model for COVID-19 mitigation Version=1.5 Date=24/03/2020			
#	Scenario	Expected Outcome	Passed
1	SIR, Naïve pop; No new agent; pImmune=0%; nStartInf=0;	No epidemic	Y
2	SIR, Naïve pop; Novel agent; Not contagious; pImmune=0%; nStartInf=1; nContacts=0; pConversion>0%;	No epidemic	Y
3	SIR, Naïve pop; Novel agent; Contagious; Not Infectious; pImmune=0%; nStartInf=1; nContacts=25; pConvert=1%; tSus2Inf=0; pInf2Rec=100%; tInfToRec=0; pInf2Sev=0%;	No epidemic	Y

As well as verifying the internal structure of the model, it is important to test the new model using an alternative modelling method and some validated test cases. For example, the expected output of specific scenarios can be predicted using analytical methods such as a proven mathematical equation.

In this context, we can use the equation for the *herd immunity threshold* (Eq. 2) to verify that if the proportion of the starting population that is already immune is below this threshold, then the SIR model predicts that an epidemic does happen; and when it is above this threshold the model predicts an epidemic does not happen.

*Figure 10. Verification test for Ro = 1.75 (HIT = 1-1/1.75 = 0.43) and pImmune=0.3 so the susceptible population is 50,000,000 * (1-0.3) = 35,000,000. The chart shows the Excel model predicts an epidemic which is the expected output. So, this verification test is passed.*

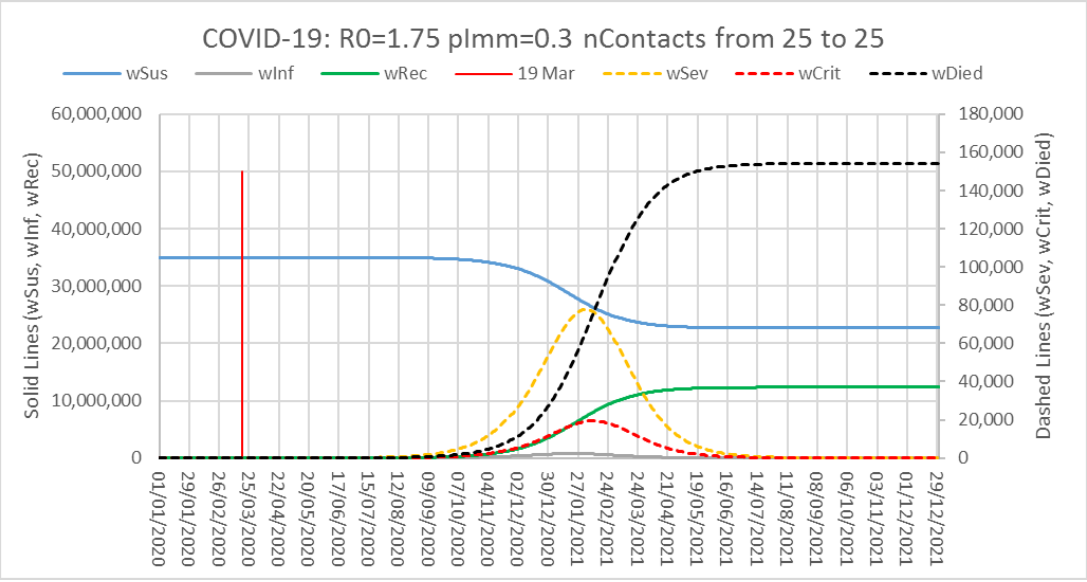

Figure 11. Verification test for Ro = 1.75 (HIT = 1-1/1.75 = 0.43) and pImmune=0.5 where the Excel model predicts no epidemic which is the expected output; so this verification test is also passed. Note that a low level of illness develops which remains stable over time and is called endemic disease.

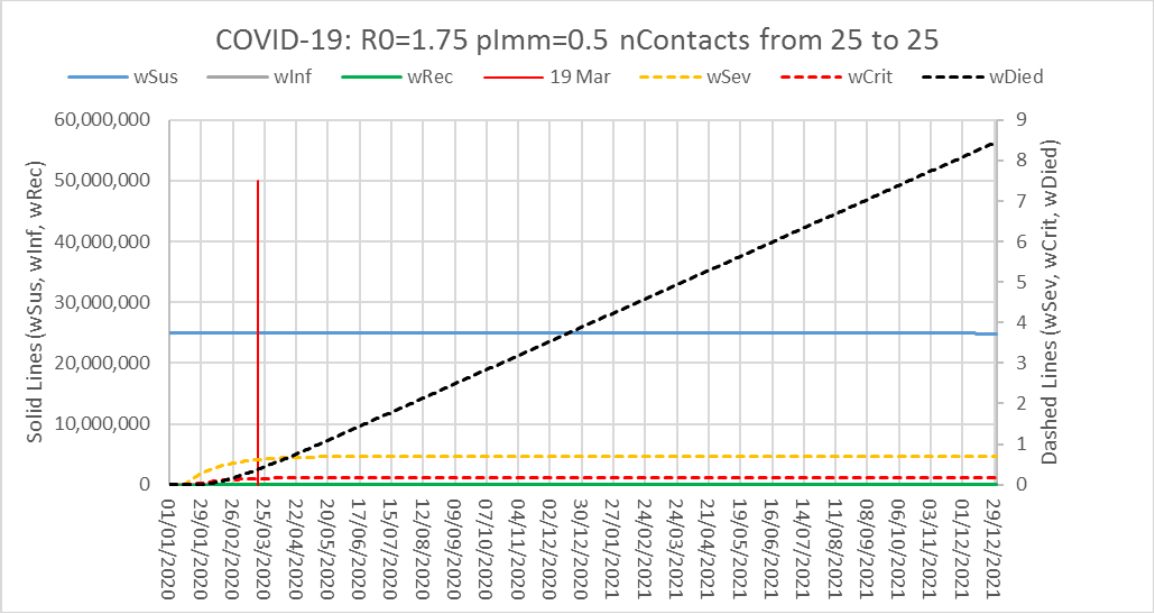

5. Implement

The first stage of implementation is to calibrate the generic model so that it matches a local context using historical data from that context. This essential calibration step allows the model inputs to be tuned to the specific system so that future predictions from the calibrated model are valid for that system. In this context the proportion of the population that has baseline immunity (pImmune) and R_0 for the remaining susceptible compartment need to be estimated using historical data.

Note: This model has not yet been calibrated so the demonstrations below are illustrative but realistic.

With a verified and calibrated model, we can define a series of experiments to answer specific questions that the model has been designed to answer by showing the predicted severity and time course of an epidemic. The two primary questions are:

Q1: Can an epidemic be supressed or flattened-and-delayed by the population adopting a social-distancing behaviour?

This question can be answered by running the model over a range values of the inputs that drive R_0 and observing the predicted system behaviour. Fig. 7 shows the stocks, flows and signals related to this question which link back to an area for the model inputs (Fig 12).

nContacts	25.0	25.0	*per day*
pConversion	1.0%	1.0%	*%*
pSus2Inf	25.0%	25.0%	*% per day*
tSus2Inf	2		*days*
pInf2Rec	100%		*%*
tInf2Rec	7		*days*

Figure 12. Screenshot of the relevant section of the input area. nContacts refers to the average number of opportunities for an individual to be exposed to COVID-19 virus in a day and pConversion refers to the probability that one exposure results in an infection in a susceptible individual. tSus2Inf represents the latent period (lag#1) in days and tInf2Rec represents the infectious period (lag#2) in days (Figs 2, 6). pInf2Rec=100% means that all infected individuals will become recovered,

*(non-infectious) so there are no persistently infectious carriers. R_0 is calculated using the formula nContacts*pConversion*tInf2Rec and represents the number of people that one infectious person infects.*

The actual values of *nContacts* and *pConversion* are not known, but R_0 and the infectious period for COVID-19 have been roughly estimated as 1.75 and 7 days respectively from the initial epidemiological studies in China. It is important to note that R_0 is context specific and may be much larger in a different context.

So, as a starting point, the estimated value for *pSusToInf* is 1.75/7 = 0.25 and the impact of social distancing can be simulated by reducing *nContacts* progressively and observing the effect on R_0 and the SIR model behaviour. In this case the latent period has been set to 2 days so this is now called a SEIR model.

Q2: How many staffed emergency and critical care beds would an average acute hospital need to make available just for patients with COVID-19 pneumonia?

COVID-19: R0=2.38 pImm=0 nContacts from 34 to 34				
Inputs		Day 83	**Units**	
dt	1.0		days	
StartDate	01/01/2020		dd/mm/yyyy	
Intervals	730			This is an uncalibrated and unvalidated model and should only be used for illustrative and eductational use. It cannot be used to make predictions.
N	50,000,000			
pImmune	0%		%	
nStartInf	1.0			
nContacts	34.0	34.0	per day	
pConversion	1.0%	1.0%	%	
pSus2Inf	34.0%	34.0%	% per day	
tSus2Inf	2		days	
pInf2Rec	100%		%	
tInf2Rec	7		days	
pRec2Sev	10%		%	
tRec2Sev	5		days	
pSev2Crit	25%		%	
tSev2Crit	4		days	
tSev2Disch	7		days	
pCrit2Death	50%		%	
tCrit2Disch	7		days	
tCrit2Death	7		days	
pDeath	1.25%		%	
tDouble	3.5		days	
Ro	2.38	2.38		

Figure 13. Screenshot of the whole input area for the model.

To predict the peak demand for acute care and critical care beds we also need to know the flow conversion rates and the touch times for the Hospital section of the stock-and-flow model (Fig. 7). The values shown here are based on the range of approximate values from the data shared by countries with more experience to date.

The inputs in blue are determined by the nature of the illness, and the specific clinical pathways and policies. The predicted bed load is driven by the time pattern of infections, which is in turn driven by the inputs shown in Fig. 12.

pRec2Sev = proportion of infected patient becoming severely ill and requiring admission.
pSev2Crit = proportion of severely ill who become critically ill and require ventilation.
pCrit2Death = proportion of critically ill patients who do not survive.

*pDeath = proportion of infected cases that die = pRec2Sev * pSev2Crit * pCrit2Death.*
tDouble = the initial doubling time of the infectious population.
R_0 is the basic reproduction number which determines epidemic potential (R_0>1.0).

Figure 14. The predicted course of the epidemic for a "do nothing" or "worst case" scenario for a susceptible population of 50,000,000 using the inputs shown in Fig. 13 and assuming there is no collateral mortality caused by other patients unable to access health care services. It predicts that, at the peak load in early July, the system would need over 700,000 acute care beds and over 180,000 critical care beds.

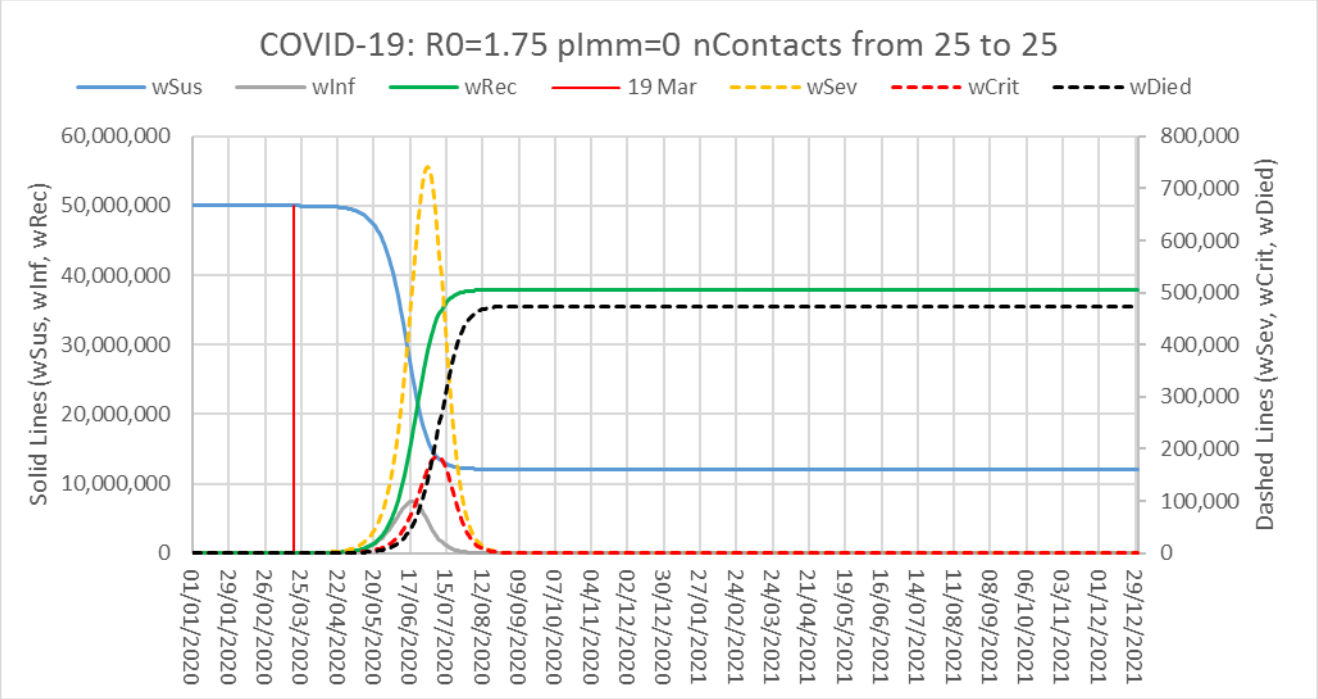

Given that we currently have no specific treatment and no vaccine for COVID-19, our only effective control lever is to reduce spread using social distancing. This implies that both questions above can be answered with one controlled experiment where the value of *nContacts* is reduced at a point in time after the start of the simulation period and the effect of the intervention observed.

Theory predicts that if R_0 can be reduced to less than 1.0 a growing epidemic can be suppressed, and this has been observed in the first countries to be affected by COVID-19 such as China, Taiwan, and South Korea who implemented "lock-down" strategies early. On 16th March 2020, the last of the temporary hospitals built in Wuhan was closed which is evidence that the COVID-19 epidemic can be supressed, and many deaths avoided. On the same date the UK government recommended an increase in the level of social distancing and has escalated the degree progressively since.

An expected consequence of supressing and epidemic is that little additional population immunity may be generated, and the remaining susceptible population remains at risk of the spontaneous development of further epidemics. That prediction assumes (a) there is low baseline population immunity and (b) recovery implies a high level of individual immunity. Currently we do not have hard evidence to validate either assumption.

Outcomes

Figure 15. Predicted effect of reducing nContacts from 25 to 22 on 19/03/2020 (equivalent to reducing R_0 from 1.75 to 1.54) and sustaining the reduction. Compared with Fig. 14 the peak demand is delayed by several weeks and reduced to 460,000 acute beds and 115,000 critical care beds.

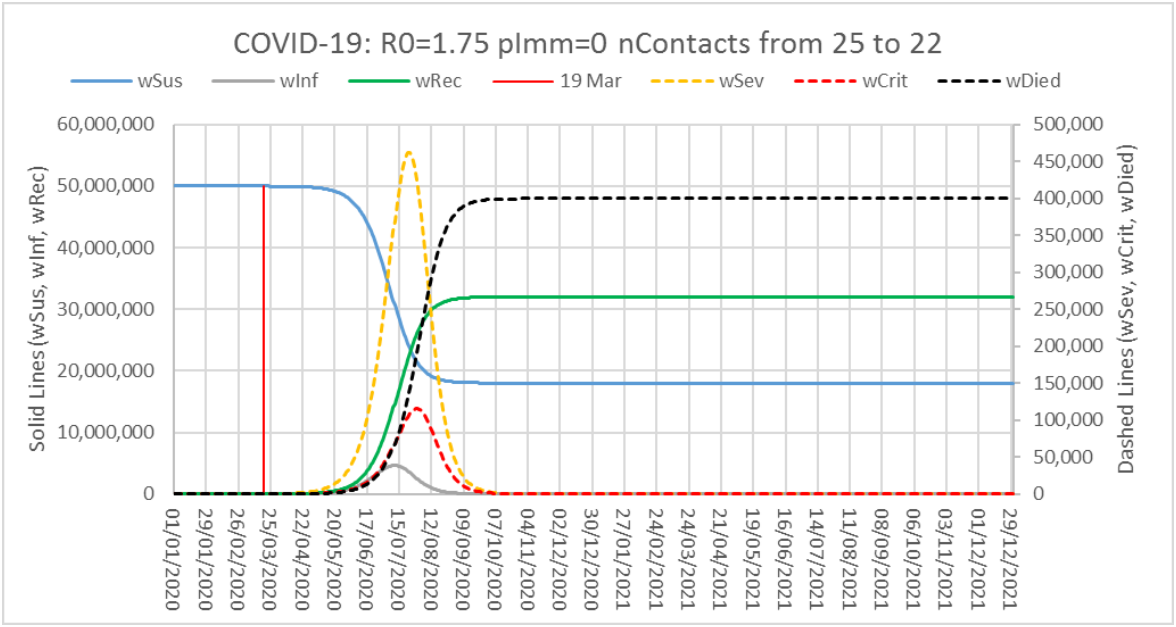

The change from Fig. 14 to Fig. 15 suggests that social distancing, if adopted universally, can have a significant life-saving effect by making it harder for the virus to spread and also generate more than enough population immunity to remove the risk of a future epidemic (HIT = 21,500,000 for R_0 of 1.75 in this example). This suggests that with a greater degree of social distancing maintained for longer, the necessary herd immunity threshold could still be achieved, and the mortality reduced still further.

Figure 16. Predicted effect of reducing nContacts from 25 to 19 on 19/03/2020 (equivalent to reducing R_0 from 1.75 to 1.33) and sustaining the reduction. Compared with Fig. 14 the peak demand is further delayed and reduced to 210,000 acute care beds and 53,000 critical care beds.

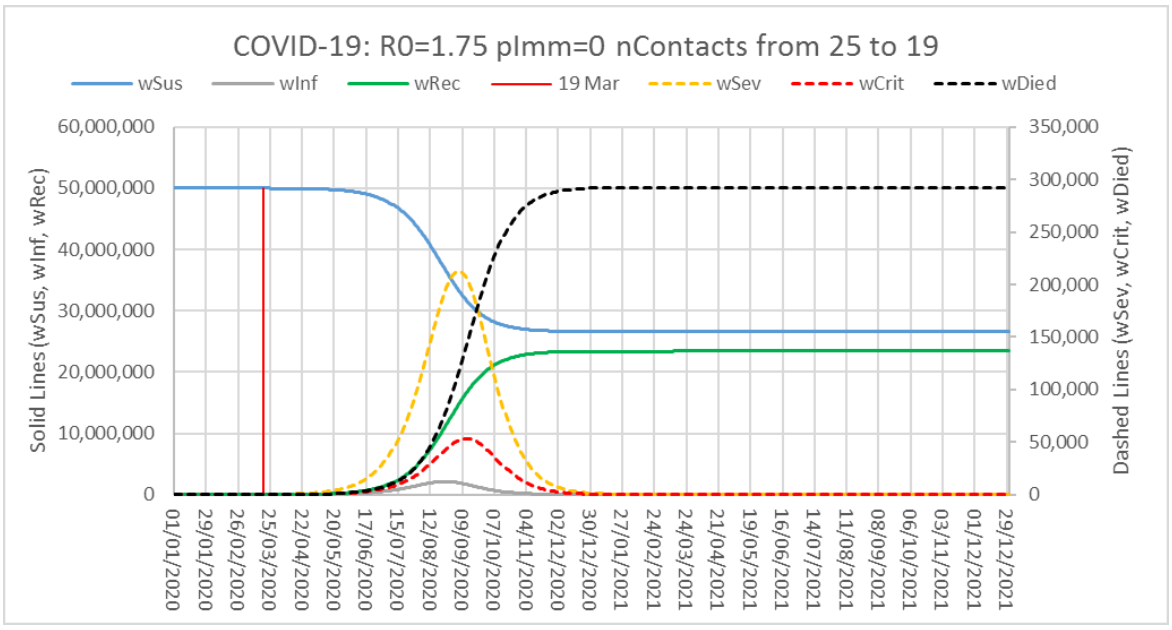

Figure 17. Predicted effect of reducing nContacts from 25 to 16 on 19/03/2020 (equivalent to reducing R_0 from 1.75 to 1.12) and sustaining the reduction. Compared with Fig. 14 the peak demand is reduced to 36,000 acute beds and 9,000 critical care beds.

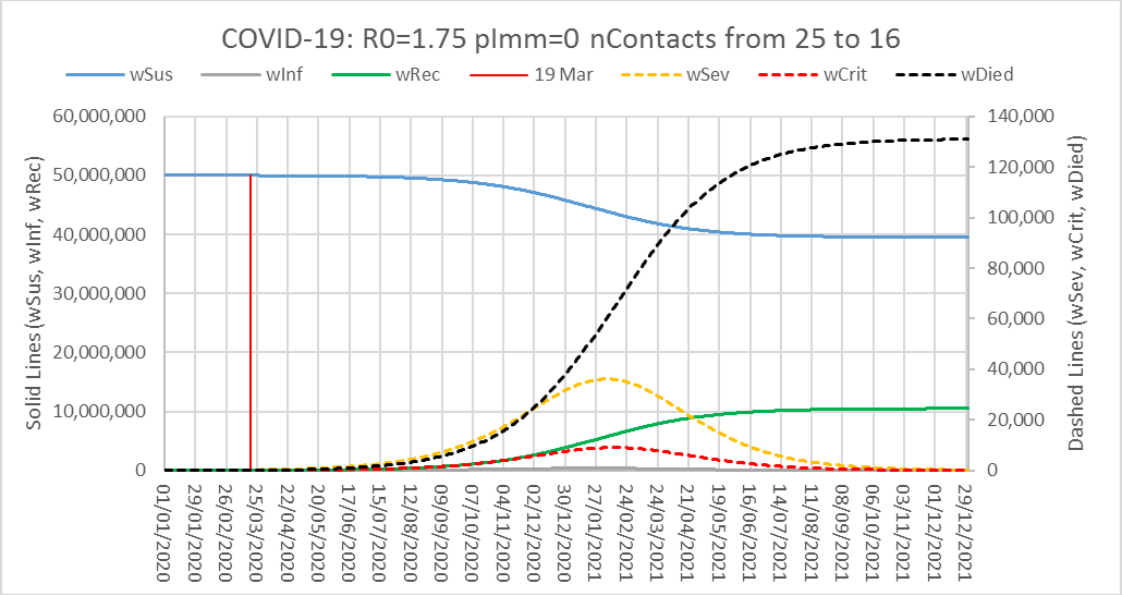

Given that we currently have no other options, reducing the spread of COVID-19 through social distancing is an effective intervention to mitigate this disease and avoid bigger humanitarian catastrophe.

Figure 18. Predicted effect of reducing nContacts from 25 to 14 on 19/03/2020 (equivalent to reducing R_0 from 1.75 to 0.98) and sustaining the reduction. Compared with Figs 14-17, the exponential growth is supressed and the delay in reducing the peak demand is only because there are already infected patients in the system.

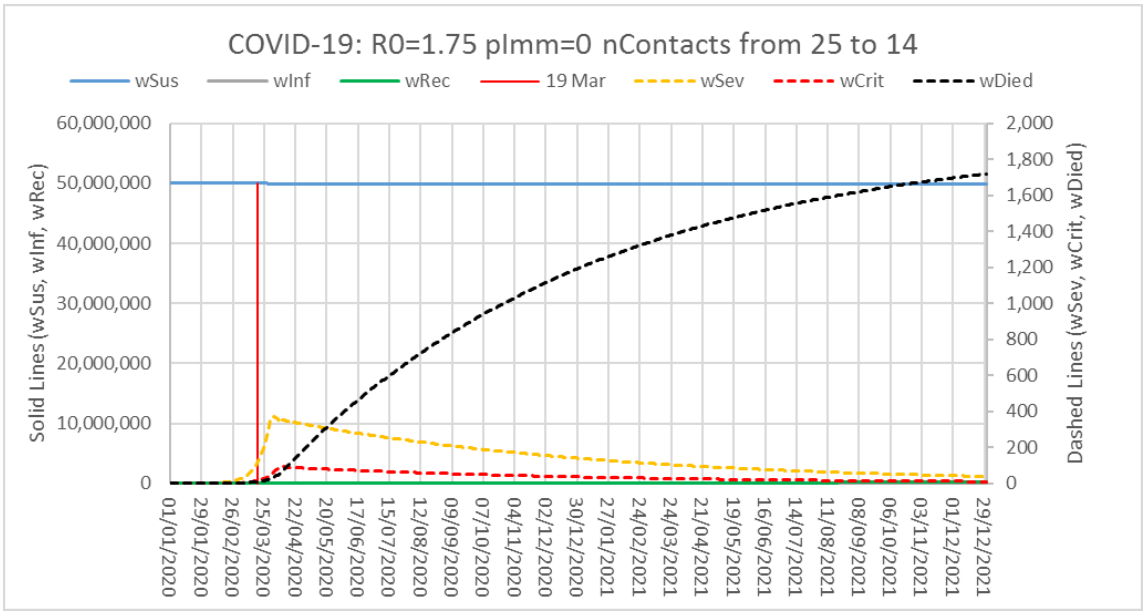

Fig. 18 confirms that with enough collective population behaviour change to limit ongoing spread of the COVID-19 virus in the community, a developing epidemic can be supressed, the acute hospitals will not be massively overwhelmed, and the mortality rate will be dramatically reduced. However, the herd immunity threshold (HIT) is not achieved so ongoing social distancing, community testing, self-reporting, self-

isolation, contact tracing and quarantine would need to be continued until the development of an effective vaccine.

6. Validate

The validation test of a predictive model is for the measured effect of an intervention to match the predicted effect. If the validation test is passed it implies that the model is fit for purpose (i.e. designed and built correctly). This model has not yet reached the validation stage, so the next step is to calibrate it using historical data and as interventions, such as nationwide social distancing, are implemented the model will be updated and re-calibrated using new data.

Reflections

It is reassuring that a realistic model of an epidemic that affects millions of people can be built using some basic information, tried-and-tested systems engineering techniques, and an Excel spreadsheet. It is hoped that sharing this insight will help inform the debate on how to effectively manage the rapidly evolving COVID-19 pandemic and reduce the human and societal cost.

There are still many unknowns such as the baseline proportion of the community who are susceptible, the context-sensitive value of R_0 and whether infection confers long-standing immunity. As the real story unfolds there will be more data to use to fine tune the model and to guide future planning. For example, it may be that the level of background immunity is much higher than zero and that the COVID-19 agent has a higher R_0. So, despite the global threat there is good reason to maintain hope that if social distancing is adopted and coordinated on a national scale, it can provide the necessary leverage to mitigate the risk.

The complex and counter-intuitive behaviour of a complex adaptive system and its sensitivity to the inputs implies that, to be able to sustain predictive power, the model needs to co-evolve alongside the system and be updated as new data comes in. This implies that the necessary knowledge and skills to use and update such models is embedded within the health care service itself. In Part 2 we will describe the model calibration phase and demonstrate a more detailed model for assisting local decision-making about the predicted demand for specific acute hospital resources.

Lessons

1. The COVID-19 epidemic can be modelled using a structurally simple system dynamics model.

2. This system dynamics model can be easily implemented using widely available spreadsheet software.

3. The illustrative model predictions match the growing body of evidence that social distancing may be an effective way to effectively manage the COVID-19 epidemic and to minimise harm.

4. As the real story unfolds the model can be fine-tuned to provide more accurate insight and guidance.

5. The importance of building capability for such modelling should be a priority for all NHS organisations.

Glossary of Terms

Asymptomatic – a disease stage where the individual does not exhibit symptoms.

Basic reproduction number (R_0) - the expected number of new cases directly generated by one case in a population where all individuals are susceptible to infection.

CAS (Complex Adaptive System) - a system in which a perfect understanding of the individual parts does not automatically convey a perfect understanding of the whole system's behaviour.

CLD (Causal Loop Diagram) – a visual representation that aids understand of how system variables are interrelated.

Critical illness – an illness that requires specialist physiological support such as ventilation.

Disease – the failure of a physiological system.

Effective reproduction number (eR) - the expected number of new cases directly generated by one case in a population where **not** all individuals are susceptible to infection.

Epidemic - a widespread occurrence of an infectious disease in a community at a specific time point.

Exposed – when an individual has encountered a disease-causing agent which is necessary for infection.

Incidence – the number of new cases of a disease in an interval of time.

Incubation period – the interval between infection and the onset of symptoms.

Infection – the invasion and multiplication of microorganisms such as bacteria, viruses, and parasites that are not normally present within the body.

Infectious – the state when and individual is infected and can transmit the pathogen to other individuals.

Immunity – an individual's resistance to infection or re-infection by a pathogen.

Latent period – the interval between the onset of infection and becoming infectious.

Pandemic – an epidemic spreading over multiple countries.

Pathogen – a disease-causing infective agent such as a virus.

Pneumonia – inflammation of the lungs caused by bacteria or virus where the air sacs are filled with pus.

Prevalence – the number of cases of a disease at a specific time point.

Prevention – lack of disease occurrence despite exposure to a pathogen.

Prophylaxis – a strategy to prevent a disease such as an immunization programme.

Recovered – the transition from an infectious state to a non-infectious state.

S&F (Stock-and-Flow) – a form of computer simulation of system behaviour.

SD (System Dynamics) – an approach to understanding the nonlinear behaviour of complex systems over time using stocks, flows, internal feedback loops, table functions and time delays.

SEIR (Susceptible-Exposed-Infectious-Recovered) – a multi-compartment model of infectious disease where the latent period is greater than zero.

Sepsis – a serious condition resulting from the presence of harmful microorganisms in the blood or other tissues and the body's response to their presence, potentially leading to the malfunctioning of various organs, shock, and death.

Serial interval - the time interval between the onset of symptoms in a primary case and the onset of symptoms in a secondary case infected by the primary case.

Severe illness – an illness that requires hospital admission for diagnosis and treatment.

SIR (Susceptible-Infectious-Recovered) – a multi-compartment model of infectious disease where the latent period is zero.

SIRS (Systemic Inflammatory Response Syndrome) – an inflammatory state affecting the whole body.

Susceptible – the state of insufficient resistance to infection or re-infection by a pathogen.

Symptomatic – a disease state where the individual does exhibit symptoms such as cough, pyrexia, etc.

Transmission – the passing of a pathogen from an infected host individual to another individual.

Virus – a small infective agent that can only multiply within the living cells of a host.

Table 1. Description of the system variables used in the Excel implementation.

wSus	Susceptible compartment stock
qSus2Inf	Susceptible to Infectious flow
wInf	Infectious compartment stock
qInf2Rec	Infectious to Recovered (non-infectious) flow
wRec	Recovered (non-infectious) stock
qSev	Flow of severe cases (admission demand)
wSev	Severe cases stock (non-ventilated)
qSev2Disch	Flow of severe to discharge (part of discharges)
qSev2Crit	Flow ofsevere to critical
wCrit	Critical case stock (ventilated)
qCrit2Disch	Flow of critical to discharge (part of discharges)
qDied	Flow of deaths
wDied	Stock of deaths (total mortality)

Note: that epidemiologists use "recovered" to mean "not-infectious" while in normal parlance we interpret it as "not ill" i.e. asymptomatic. Also, the assumption used in the model is that when a sick COVID-19 patient is admitted to hospital they become non-infectious to others in the Susceptible compartment because health care staff use infection control/isolation/personal protection to prevent spread. The patient is clearly not "recovered" in the clinical sense because they are far from asymptomatic!

References

[Chen N] Chen N, Zhou M, Dong X, et al. Epidemiological and clinical characteristics of 99 cases of 2019 novel coronavirus pneumonia in Wuhan, China: a descriptive study. *Lancet* 2020; **395:** 507–13.

[WHO 11-03-2020] World Health Organization. WHO Director-General's opening remarks at the media briefing on COVID-19 - 11 March 2020. 2020. https://www.who.int/dg/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---11-march-2020 (accessed March 11, 2020).

[PHE] Public Health England. COVID-19: investigation and initial clinical management of possible cases. 2020. https://www.gov.uk/government/publications/wuhan-novel-coronavirus-initial-investigation-of-possible-cases/investigation-and-initial-clinical-management-of-possible-cases-of-wuhan-novel-coronavirus-wn-cov-infection.

[DHSC] Department of Health and Social Care. Coronavirus action plan: a guide to what you can expect across the UK. 2020. https://www.gov.uk/government/publications/coronavirus-action-plan/coronavirus-action-plan-a-guide-to-what-you-can-expect-across-the-uk.

[Michie S] Michie S, West R, Amlôt R. Behavioural strategies for reducing covid-19 transmission in the general population. BMJ. 2020. https://blogs.bmj.com/bmj/2020/03/03/behavioural-strategies-for-reducing-covid-19-transmission-in-the-general-population/.

[Sandman PM] Sandman PM. Responding to Community Outrage: Strategies for Effective Risk Communication. 2012.

[Forrester JW] Forrester JW. Counterintuitive behavior of social systems. *Theory Decis* 1971; **2:** 109–40.

[Meadows DH] Meadows DH, Wright D. Thinking in Systems: A Primer.

[Wang 24-02-2020] Wang C, et al. A novel coronavirus outbreak of global health concern. Lancet 2020; **395;** 470-473. https://doi.org/10.1016/S0140-6736(20)30185-9 [Published online 24/01/2020].

[Qun 29-01-2020] Qun Li *et al.* Early Transmission Dynamics in Wuhan, China, of Novel Coronavirus–Infected Pneumonia. *N Engl J Med;* January 29, 2020: 1-9. https://doi.org/10.1056/NEJMoa2001316 [Published 29/01/2020]

[Wu 31-01-2020] Wu JT, Leung K, Leung GM. Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *Lancet* 2020; **395:** 689-97. https://doi.org/10.1016/S0140-6736(20)30260-9 [Published online 31/01/2020]

[Wu 03-02-2020] Wu F. *et al*. A new coronavirus associated with human respiratory disease in China. *Nature* 2020; **579:** 265-269. https://doi.org/10.1038/s41586-020-2008-3 [Published online 03/02/2020]

[Fisher 17-02-2020] Fisher D, Heymann D. Q&A: The novel coronavirus outbreak causing COVID-19. *BMC Med* 2020; 18-57. https://doi.org/10.1186/s12916-020-01533-w [Published 17/02/2020]

[Xu 19-02-2020] Xu B. Open access epidemiological data from the COVID-19 outbreak. *Lancet Infect Dis* 2020 https://doi.org/10.1016/S1473-3099(20)30119-5 [Published online 19/02/2020]

[Lin 27-02-2020] Lin Q e*t al.* A conceptual model for the coronavirus disease 2019 (COVID-19) outbreak in Wuhan, China with individual reaction and governmental action. *Int J Infect Dis* 2020; **93:**211-16. https://doi.org/10.1016/j.ijid.2020.02.058 [Accepted 27/02/2020]

[Helliwell 28-02-2020] Helliwell J *et al.* Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *Lancet Glob Health* 2020; 1-9. https://doi.org/10.1016/S2214-109X(20)30074-7 [Published online 28/02/2020]

[Qifang 04-03-2020] Qifang Bi et al. Epidemiology and Transmission of COVID-19 in Shenzhen China: Analysis of 391 cases and 1,286 of their close contacts. *MedRxIV* https://doi.org/10.1101/2020.03.03.20028423 [Posted 04/03/2020]

[Anderson 06-03-2020] Anderson RM, *et al*. How will country-based mitigation measures influence the course of the COVID-19 epidemic? *Lancet* 2020; 1-4. https://doi.org/10.1016/S0140-6736(20)30567-5 [Published online 06/03/2020]

[Ferguson 16-03-2020] Ferguson NM et al. Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. *Imperial College Covid-19 Response Team*. https://doi.org/10.25561/77482 [Published 16/03/2020].

Abstract: "The global impact of COVID-19 has been profound, and the public health threat it represents is the most serious seen in a respiratory virus since the 1918 H1N1 influenza pandemic. Here we present the results of epidemiological modelling which has informed policymaking in the UK and other countries in recent weeks. In the absence of a COVID-19 vaccine, we assess the potential role of a number of public health measures – so-called non-pharmaceutical interventions (NPIs) – aimed at reducing contact rates in the population and thereby reducing transmission of the virus. In the results presented here, we apply a previously published microsimulation model to two countries: the UK (Great Britain specifically) and the US. We conclude that the effectiveness of any one intervention in isolation is likely to be limited, requiring multiple interventions to be combined to have a substantial impact on transmission. Two fundamental strategies are possible: (a) mitigation, which focuses on slowing but not necessarily stopping epidemic spread – reducing peak healthcare demand while protecting those most at risk of severe disease from infection, and (b) suppression, which aims to reverse epidemic growth, reducing case numbers to low levels and maintaining that situation indefinitely. Each policy has major challenges. We find that that optimal mitigation policies (combining home isolation of suspect cases, home quarantine of those living in the same household as suspect cases, and social distancing of the elderly and others at most risk of severe disease) might reduce peak healthcare demand by 2/3 and deaths by half. However, the resulting mitigated epidemic would still likely result in hundreds of thousands of deaths and health systems (most notably intensive care units) being overwhelmed many times over. For countries able to achieve it, this leaves suppression as the preferred policy option. We show that in the UK and US context, suppression will minimally require a combination of social distancing of the entire population, home isolation of cases and household quarantine of their family members. This may need to be supplemented by school and university closures, though it should be recognised that such closures may have negative impacts on health systems due to increased absenteeism. The major challenge of suppression is that this type of intensive intervention package – or something equivalently effective at reducing transmission – will need to be maintained until a vaccine becomes available (potentially 18 months or more) – given that we predict that transmission will quickly rebound if interventions are relaxed. We show that intermittent social distancing – triggered by trends in disease surveillance – may allow interventions to be relaxed temporarily in relative short time windows, but measures will need to be reintroduced if or when case numbers rebound. Last, while experience in China and now South Korea show that suppression is possible in the short term, it remains to be seen whether it is possible long-term, and whether the social and economic costs of the interventions adopted thus far can be reduced".

[Wikipedia 19-03-2020] The SIR Model. [Accessed 19/03/2020] https://en.wikipedia.org/wiki/Compartmental_models_in_epidemiology#The_SIR_model

Acknowledgments

We would like to sincerely thank the community of HCSE Level 2 practitioners who offered valuable feedback on the first drafts of this initial case study and who have collectively sponsored this essay.

Authors

Mr S.R.Dodds, MA, MS, FRCS. Simon studied medicine and digital systems engineering before following a career in general and vascular surgery. In 1999, he was appointed as a consultant surgeon at Good Hope Hospital in North Birmingham and applied his skills as an engineer and clinician in the redesign of the vascular surgery clinic and the leg ulcer service. This experience led to the development of the Health Care Systems Engineering (HCSE) programme.

Dr D.T.Bradley, MB ChB, MA, MPH, PhD, MRCP(UK), FFPH. Declan is a Consultant in Public Health Medicine and Clinical Lecturer in Public Health in Queen's University, Belfast.

Sponsors

Mike Davidge, Phil Debenham, Martin Dowdall, Ken Fullerton, Chris Jones, Peter Kidd, Ellen Pirie, Delia Pudney, Kate Silvester, Lee Talbot, Lesley Wright.

Statement of Originality

Version

Reference	Version	Date	Document Owner
JOIS_2020_068	1.01	24/03/2020	simon.dodds@saasoft.co.uk