Network models for infection control through testing and contact tracing

AIMS

This project was about understanding the dynamics of testing, contact tracing and isolation (TTI). It is now broadly acknowledged that TTI (called Test & Protect in Scotland) is a key non-pharmaceutical intervention for controlling the spread of COVID-19. We set out to understand the amount of TTI that is needed in different circumstances. We additionally sought to understand the influence of the pattern of contacts among individuals in the population, or contact network, on the propagation of the disease through the population.

KEY FINDINGS

- It is possible to a limited extent to trade testing for contact tracing; however, the level of testing has a stronger effect. E.g., Increasing testing capacity by 10% has a larger impact than increasing tracing capacity by 10%.

- TTI on its own is not sufficient to control the spread of the disease unless it is done very aggressively when there are low levels of community transmission. Otherwise, other measures to reduce contact and transmissibility are required. This is true even without asymptomatic transmission of the disease; asymptomatic transmission makes this worse.

- The structure of the contact network makes a significant difference to the amount of adherence to isolation required for TTI to be effective. If everyone has more or less the same number of contacts, a lesser degree of adherence is tolerable whilst keeping the epidemic under control. If most people have few contacts but a few people have many, a much greater degree of adherence is required.

- It is productive to use a technique called Rule-Based Modelling which originates in molecular biology to model finite resources such as a particular available quantity of tests or contact tracing personnel and to couple disease transmission models with behavioural and other kinds of models.

  - Applied to surveillance testing to trigger and release lockdowns, we find that a very low level of surveillance testing produces similar overall results to a much higher level in terms of both attack rate and person-days in lockdown, but the timing of lockdown becomes less predictable.

  - We can easily incorporate simple social/behavioural models for phenomena such as wearing of face coverings or choosing to get vaccinated to understand how these interact with disease transmission. We have no concrete results here but have developed a promising methodology for investigating this class of question.

WHAT DID THE STUDY INVOLVE?

This project primarily involved contributions to the mathematical modelling literature. First, we studied TTI in well-mixed, unstructured populations. This setting does not include information about age or geography and contact between individuals in the population is assumed to be uniformly random. The reasons for this are twofold: at the time that we began there was not much detailed data available about COVID-19 in different subgroups, and we wished to understand TTI in a simplified setting to establish baseline expectations. To do this, we developed an extension to the standard "compartmental" model that we call SEIR-TTI.

We then proceeded to develop TTI modelling methodology in other more detailed settings. To validate SEIR-TTI, we implemented an individual-based model (IBM). We recovered the behaviour of SEIR-TTI using the much more computationally expensive standard IBM technique. We then implemented TTI in a network setting to investigate how contact networks affect the results. We also implemented TTI using a rule-based modelling (RBM) technique that is much more expressive than SEIR-TTI, enabling a larger class of phenomena to be represented, while at the same time being more transparent and amenable to analysis than IBM representations.

Having established this substrate of modelling capability in a variety of paradigms, we used these capabilities to ask a series of practical questions such as:

- What does the trade-off between testing and contact tracing look like?
- How is the required level of adherence to isolation affected by contact network structure?
- How do finite resources affect the efficacy of surveillance testing?

WHAT WERE THE RESULTS AND WHAT DO THEY MEAN?

Testing and contact tracing

The first set of results are about testing and contact tracing. They come from the paper, *Testing, Tracing and Isolation in Compartmental Models* by Simone Sturniolo, William Waites, Tim Colbourn, David Manheim and Jasmina Panovska-Griffiths which is to appear. The following figures from that paper capture the most important results.
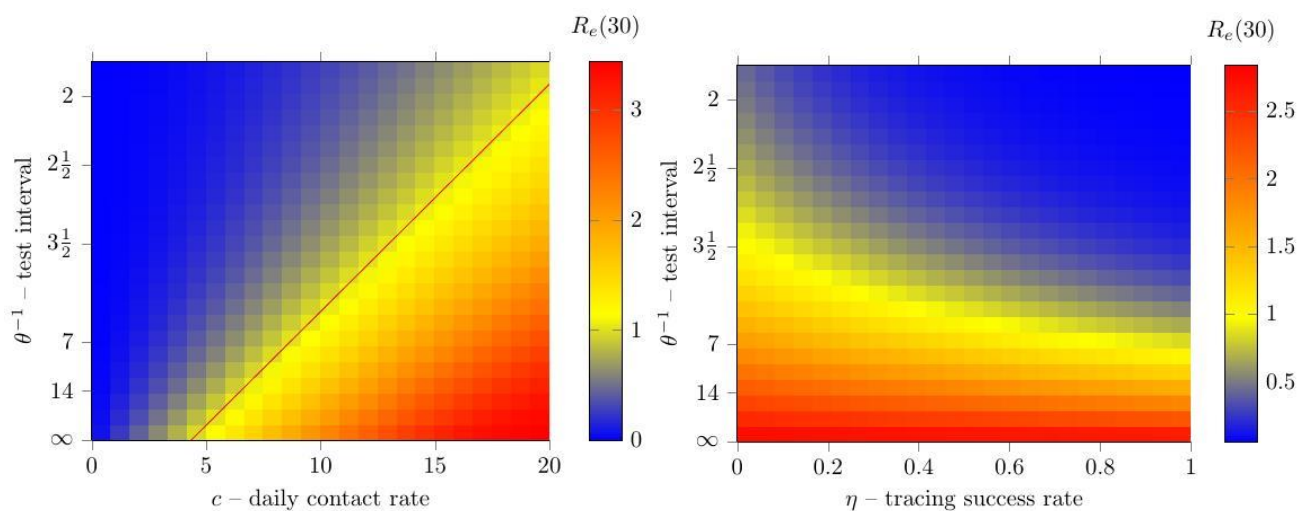
The figure on the left shows the average contact rate, how many close contacts per person per day, vs the test interval, how long (days) after becoming infected one can expect to wait for a test. The colour represents the effective reproduction number after 30 days. If it is greater than 1, the epidemic is growing, if it is less than 1 the epidemic is shrinking. The red line indicates the limit between these two situations. If the contact rate is approximately 4 or less, corresponding to a strict lockdown, the epidemic will dissipate without need for testing. As the contact rate increases, ever faster testing is required in order to maintain this situation.
The figure on the right shows the success rate of contact tracing, what proportion of contacts are traced and isolated, vs the test interval. Colour again represents R number after 30 days. This shows how, to an extent, some testing can be traded for contact tracing. Contact tracing amplifies the effectiveness of testing so, for example, an expected delay of a week for testing with very efficient tracing is equivalent to an expected delay of about 3 days with no tracing at all. Note that there is no asymptomatic transmission in this model. If we believe that 50% of cases are asymptomatic and are therefore not tested, this is equivalent to taking twice as long, on average to test.

In both cases, each pixel in the heat map represents the result of a simulation with parameters reflective of the natural history of COVID-19. Both figures have, on the vertical axis, the testing interval – the inverse of the testing rate. Towards the top is a very fast testing rate where infectious individuals are tested within not much more than a day, and towards the bottom, they are tested at a very slow rate, longer than the time-scale on which a test could conceivably produce a positive result. The colour represents the effective reproduction number some time after the beginning of the outbreak. Blue represents a suppressed epidemic and red an uncontrolled spread. Note the assumption is that adherence to isolation is perfect.

In the figure on the left, there is only testing and no contact tracing. The daily contact rate is varied. A pre-pandemic "normal" average contact rate for the UK is about 13. The 70% reduction in contact due to a lockdown is about 4. We can see that under lockdown conditions, the reduction in contact rate is sufficient to suppress the epidemic, as we indeed saw over the spring and early summer seasons in Scotland and the rest of the UK. As the contact rate increases, more testing is required. This increase appears to be linear: a doubling of contact rate requires a doubling in testing rate to achieve the same degree of control.

In the figure on the right, we fix the contact rate and vary the tracing success rate. The left-hand edge of the diagram corresponds to the left-hand diagram where the contact rate is 13. This shows that whereas we would have to maintain a very aggressive testing regime if we relied on testing alone, doing contact tracing helps amplify the effect of the testing. Following a path of constant colour through the diagram shows the extent to which more tracing can make up for less testing. The relationship is non-linear, and testing has a stronger influence than tracing. Again, this is under the assumption that individuals who have received a positive test or who are found by contact tracing do indeed isolate.

Adherence to Isolation

We investigated the effect of contact network structure on the required level of adherence to isolation in order to suppress an epidemic. We calibrated this model to be representative of the COVID-19 epidemic and explored disease transmission through explicit contact networks generated according to some standard rules. This model assumes an equal proportion of symptomatic and asymptomatic infection, and only those who are symptomatic are subject to testing. Illustrative examples are shown in the following figures.

WATTS-STROGATZ SMALL-WORLD NETWORK
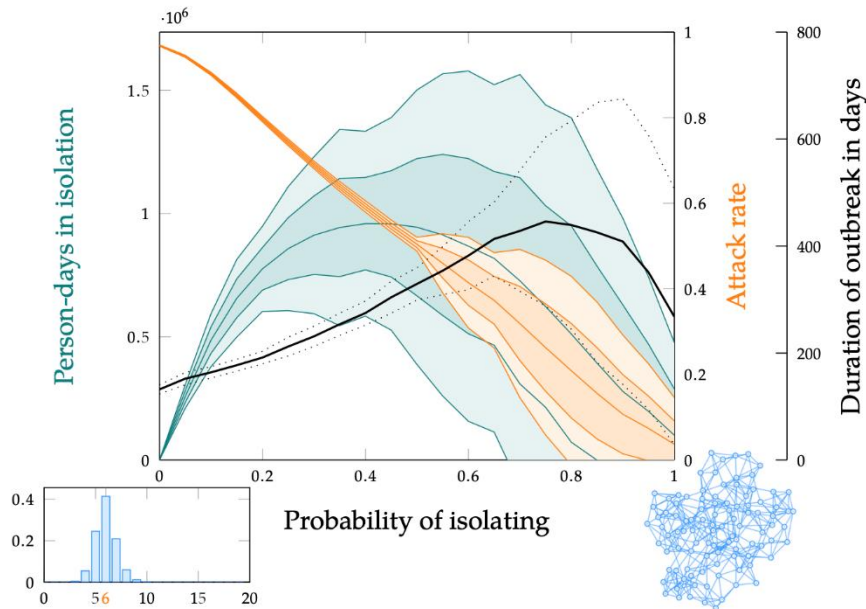
This figure shows the effect of adherence to isolation on person-days in isolation, final epidemic size or attack rate and duration for an epidemic on a network. Person-days in isolation is to be understood as a proxy for lost economic activity and final epidemic size as a proxy for mortality and long-term illness. The duration of the epidemic is self-explanatory. An indicative illustration of the structure of such a network is bottom right. The degree distribution, an important measure that says how many neighbours an individual has and can be in contact with, is shown bottom left. If there is no adherence to isolation, many people become infected and the epidemic is over quickly, nobody spends any time in isolation. If there is very good adherence to isolation, the final size of the epidemic is small and relatively less time is spent in isolation in aggregate.

Between these two scenarios we have a long epidemic with maximum time spent in isolation. There is a bifurcation at around 50% adherence to isolation where, for this kind of network, the epidemic has a chance of dying out due to local depletion of the susceptible population.

Consider first a small-world network made according to the Watts-Strogatz rule. This rule produces a network with a narrow degree distribution (bottom left of the diagram). What this means is, on average, each individual is in contact with six other individuals. The distribution is narrow so there is not much variation; some individuals may have five or seven neighbours, but nobody has fifteen or twenty. An exemplar of such a network is in the bottom right of the figure. This scenario is representative of such policies as the "rule of six" where only small gatherings are allowed consisting of a limited number of households.

The transmission model is run for plausible values of testing and tracing, and the probability of isolating is varied. The output is several different measures: person-days in isolation, the attack rate or proportion of the population that would be affected if the epidemic is allowed to run its course under those conditions, and the duration of the outbreak in days. We see that when the probability of isolation is very low, little time is spent in isolation, many people are infected, and the outbreak

is over very quickly. As the probability of isolation increases, so does the aggregate time spent in isolation and the duration of the outbreak increases. The duration of the outbreak begins decreasing again once adherence to isolation is sufficient to bring the effective reproduction number below one. Interestingly, the aggregate time spent in isolation begins decreasing at a lower level of adherence.

A particularly interesting feature of this data is that we can see that, while the orange attack rate curve decreases roughly linearly with the degree of adherence to isolation, its variance becomes quite significant after about a 50% level of adherence. The significance of this is that an increasing proportion of simulations result in the epidemic dying out. If it does not die out, the trend in the attack rate continues, but it also has an increasing chance of dying out. This means that even modest levels of adherence to isolation can have a large effect on suppression of the epidemic provided that the underlying contact network has this small-world character.
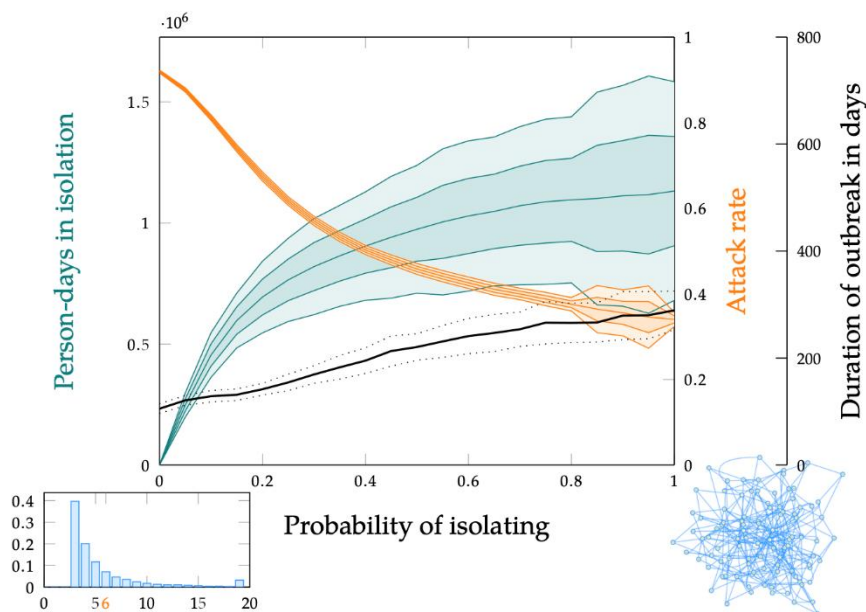
This figure differs from the previous figure in having a wider distribution of contacts as shown in the bottom left panel. While a large proportion of the population have 3 to 4 contacts there is a small number of individuals with a larger number of contacts. An example of this distribution pattern is shown in the bottom right. Increasing degrees of isolation (green lines) supresses the epidemic (attack rate in orange), but not completely as in the previous figure simulating a smaller distribution of contacts with the same mean number of contacts.

Contrasting to the above is the situation with a very different contact network structure. Here, the average number of contacts is the same, six, but the distribution is very different; it is generated according to the Barabási-Albert rule producing a scale-free network. A large proportion of the population only has three or four contacts, and there is a long tail of individuals with many contacts.

The last bucket in the degree histogram represents those individuals with twenty or greater contacts. This represents a situation where the vast majority of the population has few contacts, but a minority have many contacts. These individuals with many contacts are sometimes referred to as "super-spreaders" or individuals prone to participate in superspreading events. In these circumstances, adherence to isolation clearly improves the situation and decreases the overall attack rate, but even a very high degree of isolation is not sufficient to suppress the epidemic.

Surveillance testing with finite resources

Using a Rule-Based Model we explored the question of surveillance testing to trigger lockdowns. The general framework is introduced in the paper *Scaling up epidemiological models with rule-based modelling* by William Waites, Matteo Cavaliere, David Manheim, Jasmina Panovska-Griffiths and Vincent Danos, which is to appear. This particular application is presented in the book chapter *Mathematical modelling and data analysis as a tool for policy decision making: application to COVID-19 pandemic* by Jasmina Panovska-Griffiths, Cliff Kerr, William Waites and Robyn Stuart which is to appear.
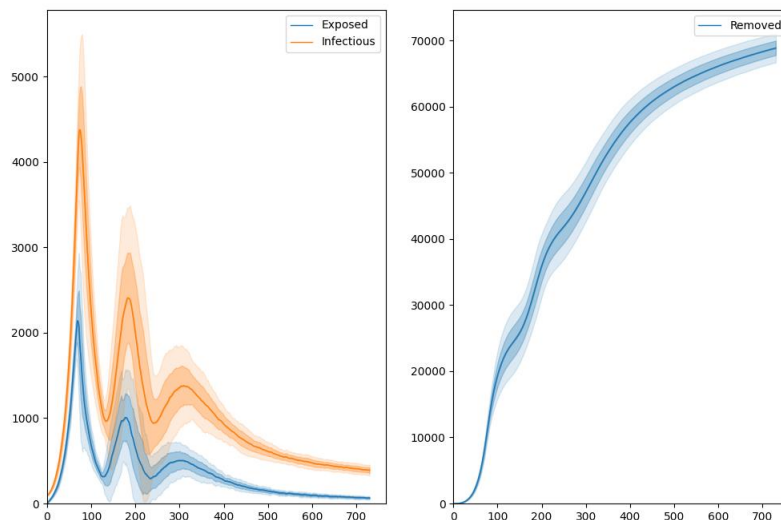

This shows an example epidemic trajectory with testing and a threshold. The horizontal axis represents the progression of the epidemic in days with the number of individuals on the vertical axis. On the left, the exposed and infectious cohort. The oscillations are due to the threshold for lock-down being met.
On the right are those who are "removed" meaning recovered or, sadly, dead. The oscillations happen predictably because tests are plentiful and so prevalence can accurately be established.

Here, we represent tests, or test kits as explicit entities. Kits are manufactured at a certain rate (in the case of the diagram above, a high rate). When an individual is tested, they become bound to a specific test, and after a short time receive a result. They can be retested after two weeks. The rate of positive tests is tracked over this window of two weeks and when it exceeds a threshold, the contact rate is drastically reduced, simulating a lockdown. When the test positivity rate falls sufficiently, the contact rate is again increased. This produces, when there are sufficient tests

available, a well-defined oscillation in the number of exposed and infectious individuals that is consistent over many simulations.

The situation is somewhat different if tests are scarce. We see no difference in the overall attack rate, the panel on the right in the figure below is very similar to the panel on the right in the figure above, but the left-hand panel below appears much less organised,
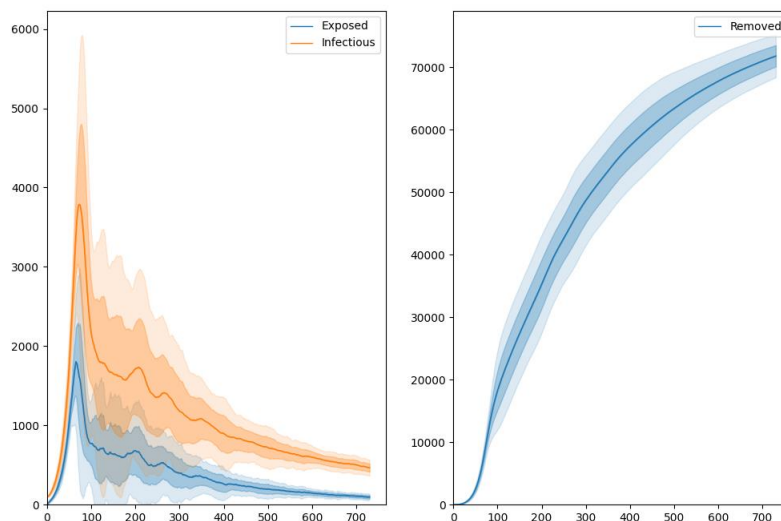
This is similar to the previous figure but in this case tests are scarce. The oscillations are no longer well defined for the average trajectory because it is no longer possible to accurately ascertain prevalence. Note that in both cases the final epidemic size, the plots on the right-hand side, are approximately equal. This shows that, for that measure, accurate estimation of prevalence does not make a large difference; what does change is the predictability of the path to that endpoint. In otherwords, it is still worthwhile to test to obtain some information, even if the signal is noisy.

The reason for this can be seen by examining an individual simulation trace, and not the mean as shown in the figures above,
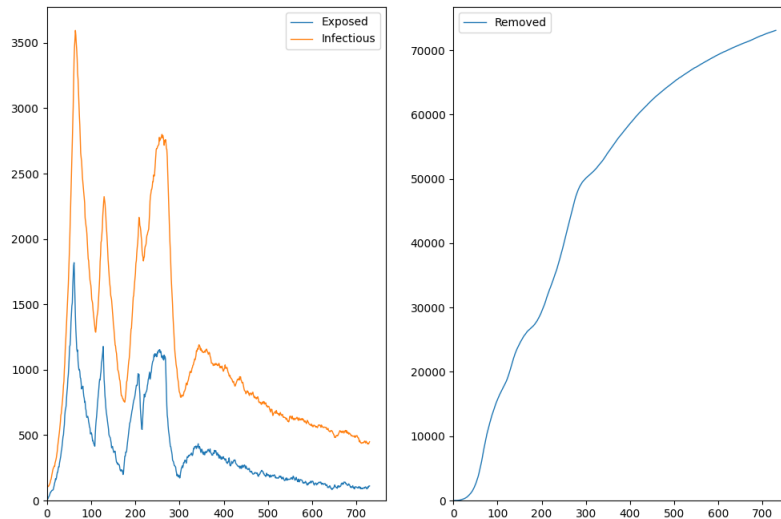
This figure shows a single trajectory from the scarce resource case. It confirms that the oscillations in infection and exposure are present but are irregular, as expected.

Here, we see irregular oscillations in the left-hand panel, and it does not appear closely related to the mean at all. In fact, it is quite different. The conclusion is that, for surveillance testing, decreasing the number of available tests by an order of magnitude makes the specific pattern of triggered interventions unpredictable but does not otherwise appreciably change the eventual situation.

WHAT IMPACT COULD THE FINDINGS HAVE?

•   Testing and contact tracing can be traded for one another to a certain extent. All else being equal, if resources are scarce, increasing investment in testing will have a greater return than increasing investment in contact tracing. This does not mean investment in testing to the exclusion of contact tracing: asymptomatic individuals who, because they will not spontaneously present themselves for testing, so a reasonable level of contact tracing is required.

- This work on TTI has formed the basis for policy recommendations and has been presented to the authorities in England as well as reported to the CSO. In particular, it formed the basis for a body of work evaluating the economic consequences of various strategies: universal testing, targeted TTI and (as a counterfactual) cycles of lockdowns. The clear conclusion is that, under conditions of high prevalence, universal testing is most effective, when prevalence is low, targeted TTI is most effective and, clearly, cycles of lockdowns are pessimal.

- The level of adherence to isolation required to contain the epidemic depends strongly on the structure of the underlying contact network. If we believe that we have a small-world network, then we can tolerate a lower level of adherence to isolation and still maintain control. If we believe that we have a preferential attachment network, with some individuals having many more contacts than others, low levels of adherence to isolation do improve the situation but very high levels are required to maintain control.

- The introduction of the rule-based formulation of epidemiological models may have long-term impact on the field, enabling concise and transparent expression of an important class of models in a way that was not previously available.

- This work has also resulted in the establishment of an infectious disease modelling collaboration between the University of Edinburgh and UCL that promises to be productive on a continuing basis.

HOW WILL THE OUTCOMES BE DISSEMINATED?

The results of this research will primarily be disseminated as academic papers in various journals. These papers are either accepted for publication or currently under review:

- Sturniolo, S., Waites, W., Colbourn, T., Manheim, D., & Panovska-Griffiths, J. (2020). *Testing, tracing and isolation in compartmental models*. doi:10.1101/2020.05.14.20101808
- Waites, W., Cavaliere, M., Manheim, D., Panovska-Griffiths, J., & Danos, V. (2020). *Scaling up epidemiological models with rule-based modelling*. arXiv: 2006.12077 [q-bio.PE]
- Panovska-Griffiths, J., Kerr, C. C., Waites, W., & Stuart, R. M. (2021). *Mathematical modelling and data analysis as a tool for policy decision making: Application to COVID-19 pandemic*. Book chapter.
- Panovska-Griffiths, J., Kerr, C., Waites, W., Stuart, R. M., Mistry, D., Foster, D., . . . Bonnell, C. (2020). *Modelling the potential impact of mask use in schools and society on COVID-19 control in the UK*.

doi:10.1101/2020.09.28.20202937

- Colbourn, T., Waites, W., Panovska-Griffiths, J., Manheim, D., Sturniolo, S., Colbourn, G., . . . Raine, R. (2020). *Modelling the Health and Economic Impacts of Population-Wide Testing, Contact Tracing and Isolation (PTTI) Strategies for COVID-19 in the UK.* doi:10.2139/ssrn.3627273

Some collaborators on this project attend government advisory panels and salient results have been communicated through those channels. There are several directions of future work that follow from this. First is the application of the modelling infrastructure to specific questions in a reactive manner. Second is the further development of rule-based modelling for infectious disease epidemiology. Though it is very productive as it is, RBM arose from molecular biology and there are some features of macroscopic populations that we would like to have are not presently convenient to express; this will be the subject of a follow-on work to extend the language. Finally, we recognise that many of the factors that influence the development of epidemics are not, strictly speaking, explicit in most epidemiological models. Behavioural and economic, for example, models are needed and these need to be combined with the disease models and not considered as exogenous phenomena. We sketched how this could be done for some simple cases with RBM but doing so thoroughly and in a multi-scale, multi-paradigm way is a crucially important, but more ambitious long-term research project.

CONCLUSION

Mathematical modelling is a useful tool for understanding the dynamics of an epidemic. This is a very different use of modelling from forecasting. We have looked at a few specific aspects of these dynamics in this work and developed a solid framework for understanding testing, tracing and isolation. We recognise that behavioural, economic and social factors have a strong influence on the course of an epidemic and should be combined with disease transmission models to understand the interactions between these coupled systems. We propose that this kind of multi-scale, multi-paradigm modelling is an important direction for future work.

Text to explain any image

RESEARCH TEAM & CONTACT

NAME or NAMES William Waites, Vincent Danos	Email address wwaites@inf.ed.ac.uk
Address IF5.38 School of Informatics University of Edinburgh 10 Crichton St Edinburgh EH8 9AB	Phone number +44 131 536-5138

ADDITIONAL INFORMATION

E.g. Date the project was completed and the amount of funding received