

Guiding Principles for Unlocking the Workforce - What Can Mathematics Tell Us?

Working Paper
May 7, 2020

This Virtual Study Group was hosted by the:
Virtual Forum for Knowledge Exchange in the Mathematical Sciences (V-KEMS)
consisting of
Knowledge Transfer Network (KTN)
Isaac Newton Institute (INI)
Newton Gateway to Mathematics
International Centre for Mathematical Sciences (ICMS)

Knowledge Transfer Network

 **Isaac Newton Institute**
for Mathematical Sciences



With support from:
a broad team gathered from across the Mathematical Sciences Community,
including representatives from the Heilbronn Institute for Mathematical Research

Report Main Author

I. David Abrahams ¹

(NM Rothschild and Sons Professor of Mathematical Sciences, Cambridge University and Director of the Isaac Newton Institute for Mathematical Sciences, Cambridge)

Contributing Authors

Demi Allen (University of Bristol), Miguel Anjos (University of Edinburgh), Ben Barber (University of Manchester), Benjamin Barrett (University of Bristol), Kirsty Bolton (University of Nottingham), Chris Budd (University of Bath), Matt Butchers (Knowledge Transfer Network), Helen Byrnes (University of Oxford), Kieran Calvert (University of Manchester), Alan Champneys (University of Bristol), Jon Chapman (University of Oxford), Edward Crane (University of Bristol), Christine Currie (University of Southampton), Sergey Dolgov (University of Bath), Christopher Doris (University of Bristol), Carina Dunlop (University of Surrey), Rosemary Dyson (University of Birmingham), Jessica Enright (University of Glasgow), Nabil Fodai (University of Nottingham), Aden Forrow (University of Oxford), Ayalvadi Ganesh (University of Bristol), Luca Giuggioli (University of Bristol), Dimitrios Gerogiorgis (University of Edinburgh), Scott Harper (University of Bristol), Rebecca Hoyle (University of Southampton), Kevin Hughes (University of Bristol), Eugenie Hunsicker (Loughborough University), Joshua Jackson (Imperial College London), Oliver Johnson (University of Bristol), Sam Johnson (University of Birmingham), Benjamin Lees (University of Bristol), Chris Martin (University of Bath), Simon Peacock (University of Manchester), Asgerdur Petursdottir (University of Bath), Tim Rogers (University of Bath), Lars Schewe (University of Edinburgh), Daniel Shiu (University of Bristol), Ben Smith (University of Manchester), Albert Sola Vilalta (University of Edinburgh), Matteo Taffetani (University of Oxford), Sam Tickle (University of Bristol), Katy Turner (University of Bristol), Arkady Wey (University of Oxford), Amy Wilson (University of Edinburgh), Helen Wilson (University College London)

¹ director@newton.ac.uk

Mathematical Principles for Unlocking the Workforce - Working Paper

WARNING: this report contains preliminary findings that have not been peer reviewed. The findings are intended to provoke further study and policy discussion and should not be treated as definitive scientific advice in response to the SARS-CoV-2 epidemic.

Whilst we expect these principles to help others formulate coherent and consistent guidelines, time has prevented any quantitative study of their effectiveness. This could be undertaken, but would require real data and time to build more detailed simulation tools. Thus, we are not able to make specific recommendations from the principles, e.g. we cannot infer that it is safe to open industry X if you follow principle Y.

Additionally, this report has been assembled in a short time frame, we have made every effort to ensure references and links are present. Where this is not the case, we apologise for the unintentional oversight. A live document here will collect any corrections and important supplementary information.

Contents

1	Introduction	5
2	Guidelines versus principles	5
3	Limitations of the Study Group	6
4	Background	8
4.1	A quick refresher on the reproduction number R , and how mathematical epidemiologists use it	8
4.2	Principles for running hot and for running cold	9
4.3	Mathematical models of contact tracing for COVID-19	10
4.4	Responsibility, non linearity and traceability	10
5	Assumption on Categories of Worker	12
6	Principles	14
6.1	The "Firebreak" Principle	14
6.2	"Fixed-Desk" Principle	16
6.3	"Table-Service" Principle	18
6.4	"Poorly-Exclusion" Principle	20
A	Two Examples of the Principles in Action	22
A.1	Basic model of the effect of fortnightly rotas, using estimated COVID-19 incubation and infectiousness parameters	22
A.2	Low degree interactions / wide isolationsphere (contact tracing)	27
B	List of Acronyms	32

1 Introduction

The purpose of this Report is to consider principles for workplace operation, based on sound mathematical arguments, that should be considered as people return to work after the UK's protective lockdown in early 2020. These principles address the issue of balancing the need to protect the population against SARS-CoV-2 (COVID19) against the economic imperative of opening up workplaces and the broader environment.

Broadly speaking this means minimising interactions, and times of interactions, to avoid infection, but doing this in such a way that a return to normality can be made as quickly as possible. However, some ways of minimising interaction are better than others in terms of their effect on

- disruption to business, workers and society
- the reproduction number and rate of decrease of COVID19 prevalence
- the probability that contact tracing and isolation fails to suppress a local outbreak that could seed a second large wave.

The Virtual Study Group (VSG) focused on applying simple mathematical models in small teams to explore these issues in different types of workplace, attempting to draw on relevant knowledge from the scientific literature. It was organised by the newly formed Virtual Forum for Knowledge Exchange in the Mathematical Sciences (V-KEMS), a joint initiative between the Isaac Newton Institute (Cambridge) the International Centre for Mathematical Sciences (Edinburgh) and the Knowledge Transfer Network.

The meeting brought together online a group of over fifty participants, all normally based in the UK. Most were academic mathematical scientists, from a range of specialities. Academics from data science, economics, epidemiology, public health, and behavioural science also took part.

2 Guidelines versus principles

The UK government must soon provide general guidelines for safe operation of workplaces that are to reopen in each stage of the release from lockdown. The guidelines themselves must be straightforward to interpret and as unambiguous as possible. They cannot be comprehensively enforced but must be largely adhered to voluntarily. Therefore public understanding of **simple principles** that underpin the detailed guidelines is crucial.

The guidelines themselves represent a choice made by the government about how best to trade off many competing needs. They will be subject to revision as evidence is gathered about their combined effect on the economy, on the functioning of society, and on the epidemic reproduction number in different parts of the UK,

The VSG was set up to respond to a request to providing "principles" about how to modify the operation of an individual workplace in order to reduce viral transmission, "all other things being equal". Note that the principles are not directly intended to enable decisions about which workplaces may reopen and which must remain shut. Instead they are intended to create guidelines that will inform the design of safe workplace operations and scheduling once it has been decided that a particular workplace is to reopen

3 Limitations of the Study Group

The Virtual Study Group (VSG) was planned and implemented in under a week, it ran for just two days, and this report was produced in one further day. Thus, the study group therefore had to produce its output without the input of complex modelling of the pandemic, modelling of social behaviour, or use of large datasets.

This report is *not intended to contribute to the scientific predictive modelling of the COVID19 epidemic.*

We mainly steered clear of the following ethical minefields:

- different rules for different people based on what a mathematical model says is their relative likelihood of being infected
- giving certain extra rights and responsibilities to holders of immunity certificates
- rules putting people at more or less risk at work depending on their domestic circumstances

While there are strong mathematical arguments in favour of each of these, their acceptability or employment are not mathematical questions.

Transmission via contaminated surfaces appears to be important for COVID19, but for reasons of time and lack of data, most work in the study group considered only direct transmission from close human contact.

The report is intended to be useful for non-experts. But having been produced collaboratively in a very short time, inevitably technical and specialist language remains.

Mathematical Principles for Unlocking the Workforce - Working Paper

There is some overlap between the principles proposed by different teams at the study group.

This report **has not been peer-reviewed**. It contains the rapid opinions of many contributors based on numerous simple mathematical models, some of which could form the basis for further scientific study.

4 Background

4.1 A quick refresher on the reproduction number R , and how mathematical epidemiologists use it

There has been much public discourse about the reproduction number R of the COVID-19 virus. We are often told that "we must reduce R " or "we must keep R below 1". The reproduction number varies over time as an epidemic proceeds, so it is sometimes denoted R_t . Its initial value near the beginning of the epidemic, in the absence of any control measures or immunity, is denoted R_0 . Thus we do not speak about trying to change R_0 , only about trying to change R . But what precisely is R ? It is the mean number of people to whom a "typical infected individual" transmits the viral infection. Notice that this is not the same as the mean number of people to whom a person drawn uniformly at random from the population would transmit the infection. Nor is it the mean number of people that "you" (being very careful and responsible) would expect to transmit the virus to, were you to be so unlucky as to acquire the virus today. Why are these not the same? Because individual characteristics and circumstances make some people very much more likely to be infected than others. This fact is very important and underlies all serious mathematical thinking about the spread of epidemics through populations. It also underlies a lot of the thinking that was done at the study group.

Whether R is greater than one ($R > 1$) or less than one ($R < 1$) is very important. When $R > 1$ then the epidemic will grow in the long run (until acquired immunity or control measures reduce R below 1). But if $R < 1$ then the epidemic will shrink and die out in the long run. This can be made into a rigorous mathematical theorem about some fairly realistic mathematical epidemic models. More importantly for studying real-life epidemics, this fact about the threshold at $R = 1$ is robust to the differences between the abstract mathematical models and real life.

When $R > 1$, knowing the value of R alone does not tell us how quickly the number of infected people will grow (usually measured by an "exponential growth rate" r , a small number such as 0.14). For this we also need to know how likely the virus is to be transmitted at different stages of the course of one typical infected individual's infection. This in turn is affected by how severely ill they become and whether they isolate at some point during the course of the infection, so it is not purely a matter of biology. Likewise when $R < 1$ then knowing the value of R does not tell us how quickly the number of infected people will shrink, (measured again by r , now a negative number). But it will eventually shrink to 0.

R also varies from place to place. It is no use if R is below 1 everywhere across the country except in one city where it is above 2 for several weeks, because that city will then experience a large outbreak, which in turn will export lots of infections to those areas where R was below 1 unless it is sealed off.

4.2 Principles for **running hot** and for **running cold**

We will use colours to emphasize the principles that are useful in two different national situations.

- **Red principles** are for pushing down R in a "running hot" situation where the proportion of infected individuals in the population remains unacceptably high and we must maintain $R < 1$ everywhere in order to stay on the downslope of the epidemic.
- **Blue principles** are appropriate for the "running cold" situation when the proportion of infected individuals is acceptably low and we may resume (close to) normal life, but dangerous outbreaks are still possible (for example seeded by imported cases). Outbreaks are suppressed by tracing the contacts of suspected or confirmed cases. The contacts are isolated either for 14 days or until a test of the suspected case is returned negative.

In the running cold situation the value of R is kept artificially low by the contact tracers, because the "typical infected individual" is likely to be found and isolated quickly before spending much time in the highly infectious state. The running cold situation is highly desirable to reach because most of normal life can resume, and we merely have to adapt our practices to "minimize the probability that an outbreak occurs that exceeds the capacity of the contact tracers to chase it down". That event would effectively return us to running hot. Running cold with contact tracing means that we can live our life in a way that would result in R much greater than 1 if it were not for the contact tracers. Running cold will not eradicate the virus because we will continue to import infectious but asymptomatic cases.

At the time of writing, May 1st 2020, the UK is still **running hot**, while South Korea is already **running cold**, after being an early centre of the epidemic.

4.3 Mathematical models of contact tracing for COVID-19

There are several recent mathematical epidemiology papers about the feasibility of suppressing COVID-19 outbreaks using contact tracing and isolation, among which we mention two. Hellewell *et al* (Feb. 2020)² studies a mathematical model of an epidemic outbreak fought with contact tracing and isolation of contacts in order to look at the relationship between R_0 (here standing for what R would be without the contact tracing and isolation, but perhaps not in the complete absence of other control measures), incubation and infectious period, probability that each contact is traced, and the time delay in tracing and isolated contacts. When $R_0 = 1.5$, they found only 50% of contacts had to be traced to contain an outbreak; for $R_0 = 2.5$ this rose to 70% and for $R_0 = 3.5$ it rose to 90%. Minimising the delay between symptom onset and isolation was also found to be crucial.

Ferretti *et al* (Mar. 2020)³ used the latest estimates then available of COVID19 infection dynamics and case count data from the Hubei outbreak to study the effectiveness of contact tracing, and concluded that traditional "manual" contact tracing would be insufficient to prevent outbreaks "in the absence of other measures to reduce transmission", but that contact tracing assisted by location-sensing and proximity-sensing technology would be required.

4.4 Responsibility, non linearity and traceability

There is a huge variety of workplaces in the UK, and it will not be possible for the UK government to provide fine-grained guidelines that apply meaningfully to each one. Instead some "responsibility" must be given to employers to use the principles to develop their own detailed COVID19 plan with a view to the safety of all users of the environment over which they have control, including their employees.⁴

There is an important but more subtle responsibility that goes beyond common sense, because it arises from a "highly nonlinear effect". Unlike ordinary health and safety rules, these new decisions to be made about a workplace's operation will have a significant effect on the risk to the population as a whole. This is because people who acquire COVID-19 on the premises may go on to seed a cascade of subsequent infections. If enough people acquire the virus in one place in a short interval of time and are not then found by the contact tracers, a local outbreak may quickly arise that could

² [https://www.thelancet.com/pdfs/journals/langlo/PIIS2214-109X\(20\)30074-7.pdf](https://www.thelancet.com/pdfs/journals/langlo/PIIS2214-109X(20)30074-7.pdf)

³ (<https://science.sciencemag.org/content/early/2020/04/09/science.abb6936>)

⁴ For information as to how the New Zealand government is asking employers to make such plans, see <https://worksafe.govt.nz/managing-health-and-safety/novel-coronavirus-covid/covid-19-safety-plan-what-you-need-to-think-about/>.

exceed the capacity of contact tracing to suppress. Such an outbreak would require a new lockdown to be imposed on the affected region.

Fortunately, contact tracing and isolation increase the direct incentive for employers to modify the pattern of human interaction on their premises in beneficial ways, because they will wish to minimise the risk of many staff being placed in isolation simultaneously due to a single chain of infection occurring on site. Employers could therefore benefit from understanding the principles that we will describe.

The feasibility of running cold by contact tracing and isolation depends on many factors, including

- the number of skilled contact tracers available⁵
- people's honesty and co-operation in admitting that they are unwell with possible COVID19 symptoms, immediately isolating themselves, and engaging with the contact tracers
- the population's willingness to be tracked, e.g. by a contract-tracing app, and
- a typical infected person's daily number of "untraceable contacts"

It is the last one of these factors that is the target of the blue principles in this report.

⁵ see <https://www.cdc.gov/coronavirus/2019-ncov/php/principles-contact-tracing.html> to understand the skill set required.

5 Assumption on Categories of Worker

There is evidence that the probability of infection in a close contact between individuals depends strongly on the length of the contact, rising rapidly after a critical time t_c of around ten minutes, say. The closeness of contact is also very important, which is the reason for imposing a critical length, l_c , typically taken as 2 metres, as a social-distancing control measure. Jobs may be arranged in a rough order that combines the duration of contacts and closeness of those contacts to quantify this infection risk per contact.

The number of contacts is also important. The average rate of infections transmitted from an infectious worker increases with the size of the group the worker interacts closely with. We have discussed, above, the difficulty caused by untraceable contacts, in the running cold setting. We combine these factors roughly into a scale going from a "closed" workplace (where a worker typically interacts with few, known co-workers) to an "open" workplace (interacting with many unknown transient co-workers or customers). The closed workplace has the advantage that the number of people that can be reached in two or three generations of the infection from a single worker is much smaller than in the open workplace situation.

These two "dimensions" seemed a good starting point for clustering different workplaces to which similar general safety principles might apply.

Accordingly, the workshop considered the four cases of long and short interaction times with open and closed groups. In each case the workshop considered principles of safe operation appropriate to each of these cases. These are illustrated in the diagram below, in which we show some representative personae in each quadrant. By focusing on each persona and the environment in which they work, it is possible both to consider principles relevant to that persona, to test them and to consider caveats for using them.

After a session on framing the problem, the study group divided into four teams, one for each quadrant of the diagram, corresponding to four categories of worker categorised according to the nature of their working-time interactions:

- Long interactions within closed communities
- Long interactions with an open community
- Short interactions within closed communities
- Short interactions with an open community.

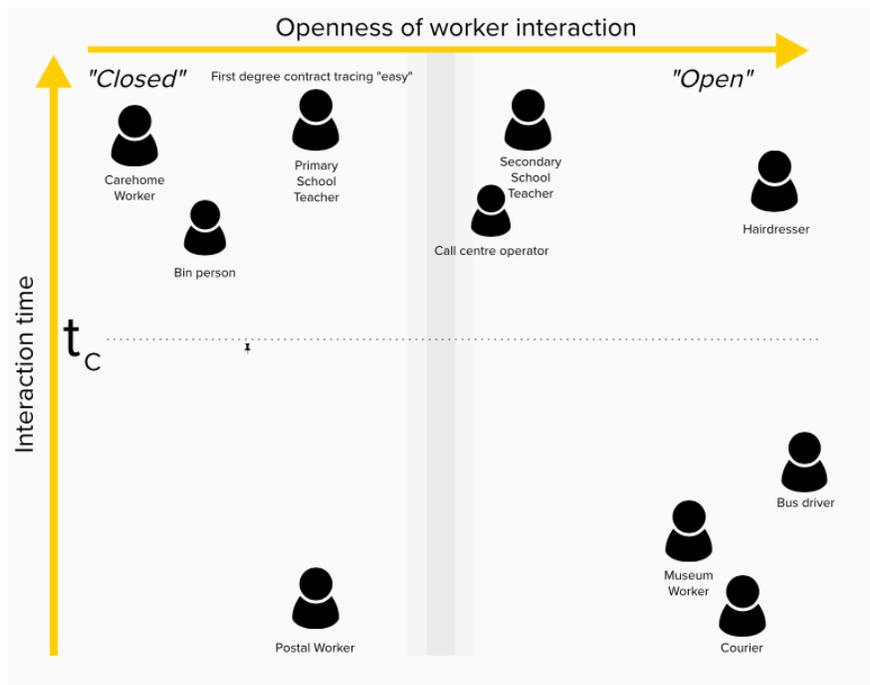


Figure 1: Framework by which the mathematical scientists approached this challenge.

This VSG Report considers 21 principles appropriate to these four different kinds of working environments. In each case we give a mathematical rationale behind the principle and test it out on appropriate simple models, undertaking simulations as appropriate. However, the work to-date has been carried out under much uncertainty about the modes of infection (both short and long term) and consequently caveats for each of the principles must be considered.

6 Principles

A total of 21 such principles were proposed by the VSG, which have been worked through in various levels of detail. We have loosely grouped these by type, but time has not allowed us to avoid unnecessary overlap. Most of these principles are listed below, together with a short description for each. At the end of this document we offer mathematical descriptions for two of the principles, as exemplars.

Note that the principles are not directly intended to enable decisions about which workplaces may reopen and which must remain shut. Instead they are intended to create guidelines that will inform the design of safe workplace operations and scheduling once it has been decided that a particular workplace is to reopen.

6.1 The "Firebreak" Principle

Partition the population / workforce as far as possible into groups that don't interact / intersect. (This is the firebreak principle used e.g. for building safety - infection might spread within the group but cannot be transmitted to other groups.)

- **"Minimising Total Contacts" Principle:** In a closed workplace such as an office where we can control exactly who interacts with whom, and where there is a policy of "if person X gets infected, everybody in their immediate neighbourhood/group is isolated", then we can limit the chances of an outbreak infecting an entire cohort by reducing the number of other people each person interacts with. The mathematical analysis for this principle is demonstrated in Appendix A.2.

- **"Islands and Bridges" Principle:**

Even when the Firebreak Principle of complete isolation cannot be fully implemented, there is still value in trying to structure workplaces into smaller groups "islands" where there are many interactions within groups, but few interactions ("bridges") between members of two different groups (see Figure 2).

These group divisions can be naturally arising from the setting (e.g. science students at a sixth form college) or artificial (e.g. ask Uber to pair clients and drivers whose surnames begin with the same letter whenever possible).

- **"Partition by Surname" Principle:** The effects of applying the above principles in a particular setting (e.g. a school or workplace) will be to some extent negated if the "islands" are mixed outside of that setting.

To minimise this, the use of labels that can be consistently applied across all settings is recommended. Automatic partitioning could be undertaken as a national

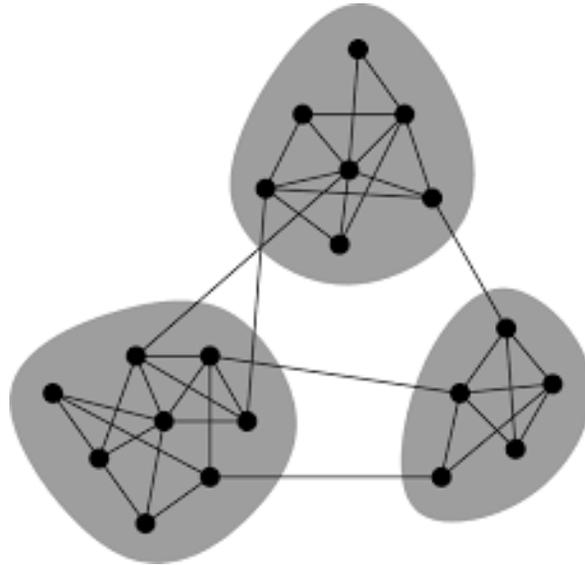


Figure 2: Grey shaded areas represent "islands" links between islands are "bridges", links within islands are not.

scheme based on a phone app, with labels created uniquely by, say, National Insurance number.

- **"Stay in two postcodes" Principle:** Encourage people and businesses not to travel or operate beyond their local area wherever possible. The oft-mentioned exponential growth of an outbreak is not possible if it is constrained in space.

We propose to organise this by postcode area, for example allowing people to move freely only in destinations within at most two postcode areas (i.e. the first three digits). This has the advantage of being easily understood and also adaptive to population density (e.g. cities have smaller postcode areas.)

- **"No Stigma" Principle:** Making behavioural decisions based on your own estimate of your risk of currently being infectious but asymptomatic could become built into everyday behaviour and become automatic and a source of pride. A kind of *smart rationing* of time allowed in crowded places?

For an incentive (e.g. a trip to the pub or hairdresser) you could voluntarily make known to everyone else what your relative probability of being infectious is. A smartphone app could score your current risk based on your recent history. Then, if necessary, you could decide (or be forced) to avoid contacts, face-to-face meetings etc for a while when your score gets above a threshold. This could be applied to both the workplace and elsewhere.

6.2 "Fixed-Desk" Principle

In the workplace maintain one-to-one or one-to-few interactions wherever possible. (This is the **fixed-desk principle** – ensure shifts times of staff are in phase; bus drivers are assigned on the same routes at the same times and with the same buses; hot-desking is banned.)

- **"Nesting" Principle:** When workers interact with a closed group of clients, the group of clients that any two different workers interact with should be distinct. If this is not possible, then one should be fully contained in the other. This may be achievable by appropriate scheduling of the workers.

Figure 3 indicates, on the left, the situation when the same group of clients is only ever served by one worker from each shift. If this is not possible, then, as shown on the right, keep clients distinct for at least one shift of the workers (here the yellow shift).

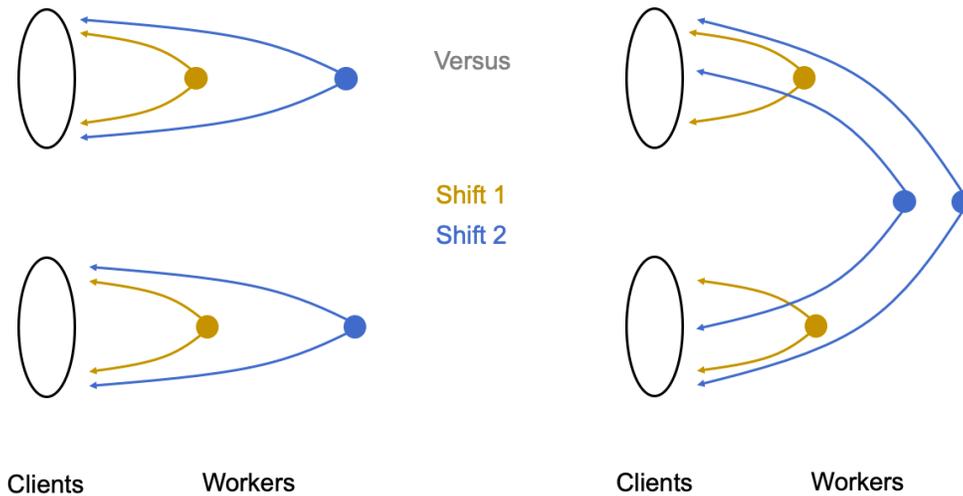


Figure 3: Nesting Principle

- **"Align Your Cohorts" Principle:** This is an extension of the nesting principle. Whenever possible people should try to align their contact networks across different activities (work, sport, school, etc). This also applies to families. This means people who work together should also work out together, their children should play together etc.

- **"Bubble Scheduling" Principle:** The principle is to separate all people in a given context (e.g. workplace) into groups (islands/bubbles), and then schedule these groups together throughout their time in that context. The aim is to create the best possible schedule for separating distinct bubbles in both space or time (see Figure 4).

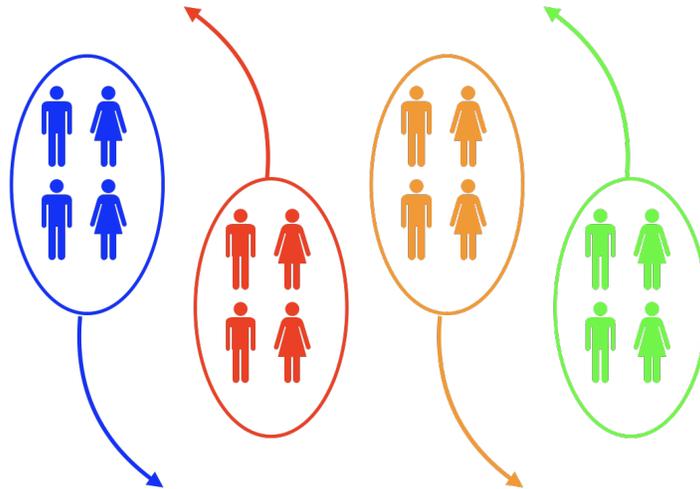


Figure 4: Scheduling the 'bubbles'.

- **"Non-mixing" Principle:** Aim to avoid mixing in corridors and shared spaces by: staggered start, finish and break times; introducing one-way systems; keeping groups in one room or "home space" for as much time as possible.
- **"Minimising Changeovers" Principle:** Schedule rotas in such a way as to minimise changeovers between groups using the same spaces. In schools this could relate to block timetabling (a bubble of students do a week of science in the science lab, followed by a week of computing in a computing lab, followed by a week of English in a standard classroom).

6.3 "Table-Service" Principle

Turn many-to-many into one-to-many interactions where possible. (This is the **table-service principle** – avoid counter service in restaurants; checkout tills could be assigned to shoppers by surname.)

- **"Manage the crowd" Principle:** In a situation where there is a large crowd with many possible interactions (such as a supermarket) manage the crowd in such a way as to minimise these interactions.
- **"Pruning" Principle:** Business in the service sector, such as at-table service in a restaurant, often involves key individuals interacting with an *open* group. (An *open* group is one with a changing population, such as shoppers in a supermarket, whereas a *closed* group has a fixed cohort.) The interactions of these should be 'pruned' so that the key individuals do not pass infection between separated groups, they spend minimal time with each group, and they have as few interactions as possible to deliver an effective service.
- **"Managing the Queue" Principle:** In certain open situations the crowd becomes a queue. An example would be a take-away or a coffee shop or a reception desk at a hotel or business. In this case the queue can be managed more easily to reduce the risk of infection. However, there are economic considerations as to how long a customer should wait in a queue, how many servers there should be, and how the safety of the queue (where people will be stationary) is affected by the use of Personal Protective Equipment (PPE).
- **"Make a Chain" Principle:** Consider a workplace in which the workers have a small number of interactions, but some are unavoidable. How can such interactions be made more safe?

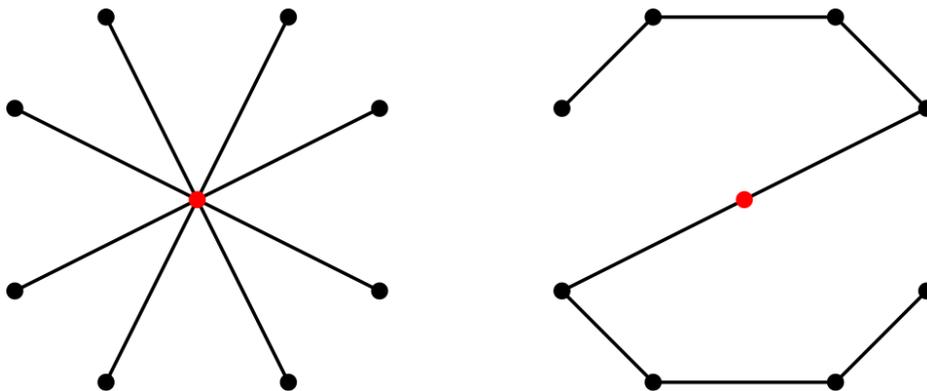


Figure 5: Making a chain between workers to increase their average-'distance'.

A configuration that appears safe is the *star graph* – everyone in the workforce interacts only with a single central point (maybe a supervisor or manager). This naively seems good because almost everyone interacts with only one person; however globally it is bad as the *average distance* (measured in some *abstract* rather than necessarily physical sense) between workers is low.

A safer model is a *chain* (see Figure 5). Rather than everyone interacting with one central node, workers interact only with their neighbours; there is no single point of failure and the average distance between workers has increased significantly. For restaurants this implies partitioning wait-staff to allow each to serve unique clusters of tables.

6.4 "Poorly-Exclusion" Principle

Adjust working/shift patterns to allow regular testing/periodic lockdown for people who have to interact with many individuals or groups. (This is the 'poorly'-exclusion principle – remove nexus staff (potential super-spreaders) before they become infectors.)

- **"Week On - Week Off Scheduling" Principle:** The key idea is that week on - week off schedules reduce transmission rates (see Figure 6). Segmenting the workforce (pruning) can also be achieved by scheduling the workforce (with days off) to remove overlap times between two teams: those who work odd weeks and those who work even weeks. This fortnightly week on - week off rota seems very practical (more practical than 4 days on - 10 days off⁶). It has the benefit that a worker who becomes infected at work is likely to spend the majority of the dangerous time when they are infectious but not yet symptomatic at home, not passing on the infection to other workers. This principle is discussed in mathematical detail in Appendix A.1.

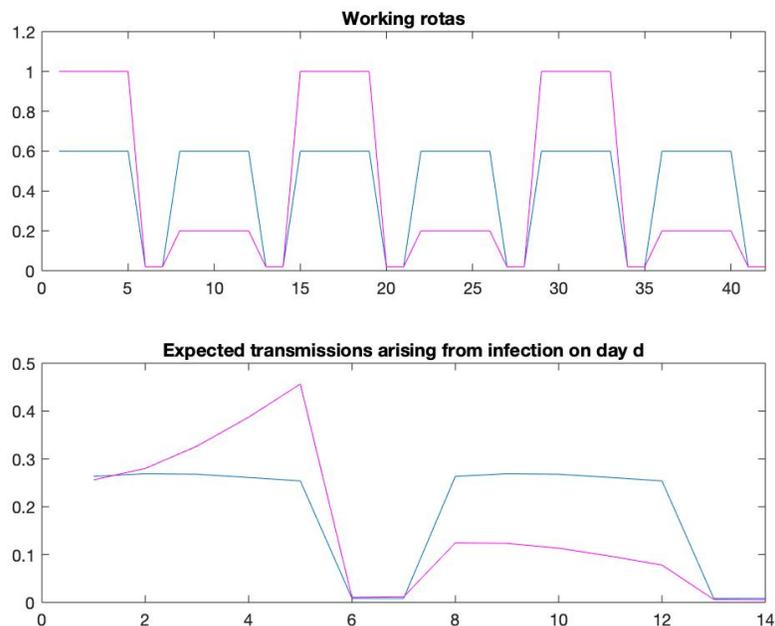


Figure 6: Simulated transmission for week on - week off working rotas. (Top) Rota profile and (Bottom) expected transmission for two rota scenarios.

⁶ Alon et al (2020+), "Adaptive cyclic exit strategies to suppress COVID-19 and allow economic activity".

- **"Traceability" Principle:** One could change business models, particularly regarding booking and payment methods to make contacts more traceable. For example, booked taxis (or ride-hailing apps) lead to a trace which contact tracers could use to connect the infected passenger to the driver and subsequent passengers. Ordinary taxis picked up at ranks or hailed on the street, with cash payment, leave no such trace. Many other informal businesses taking cash have similar problems, such as walk-in barbers.
- **"Optimize Face-to-Face Time" Principle:** Prioritise the use of face-to-face time to ensure that it is used in the best way to achieve the objectives of the organisation involved.

For schools this may be provision of PE, technology lessons, drama, science labs. For other subjects, look at a mixed approach where one lesson a week/fortnight is face-to-face and others are online. Some children could have more lessons in school than others if this is needed (e.g. vulnerable categories, poor internet access, key worker parents).

For healthcare workers, this could involve continuing / increased use of telemedicine for check-ups with patients only attending clinics when a physical procedure or examination is really necessary.

A Two Examples of the Principles in Action

A.1 Basic model of the effect of fortnightly rotas, using estimated COVID-19 incubation and infectiousness parameters

“Week On - Week Off” Principle

We use a simple model of incubation period and infectiousness of Sars-CoV-2 (a slightly simplified version of the mathematical model taken from Imperial College report 16) to understand the effectiveness of on-off work rotas with a 14-day period.

Scenario 1: A business employs two part-time workers who share a job. They can both work 5 weekdays each fortnight. Is it better for them to alternate working days or each to work one week on, one week off, alternating weeks between them?

Scenario 2: A business employs two workers, each of whom can do the other's job, or easily be trained to do so; one job is "high contact, open, high risk" and the other job is "low contact, closed, low risk". They both work Monday to Friday. What is the benefit of making them swap jobs every week?

Simple model: start with an uninfected and susceptible individual. Assume that the daily hazard rate for getting the infection while working is p_{work} and while at home is p_{home} . We will mainly look at the case where p_{home} is much smaller than p_{work} . Assume that the incubation period is random with a certain distribution (mean 5 days, variance 5 days²?). Assume that after the incubation period the infected worker becomes infectious. They then transmit infection while at work according to a Poisson process whose rate is the product of p_{work} and a time-varying infectiousness parameter, which follows a certain temporal pattern (in the Imperial College report it is the density of a Weibull distribution, but we could use something simpler). While at home they transmit at p_{home} times infectiousness. Then it is easy to compute the expected total number of infections that the worker transmits to other people over the course of a few months. (Our model for infectiousness only needs to specify the mean infectiousness on each particular day after infection occurs). We should make sure that the overall scalar constant is such that the sum over all infected days of mean infectiousness is a sensible quantity, e.g. 5, so that for the high contact - low contact rota, we just have three states instead of two: p_{high} , p_{low} , p_{home} .

Questions:

1. 5 days on / 5 days off sounds a lot more practical for scheduling workers than 4 days on / 10 days off. How much less effective is it?

2. Instead of on-off, we could do a rotation between high-contact and low-contact roles, again in a 5 days high / 5 days low pattern.
3. What proportion of people actually have the ability to do both a high-contact and a low-contact role? Maybe this is not common enough to be worthwhile making into a principle?)

Details of our model

1. discrete time (one day per discrete time), periodic on a two-week period.
2. incubation period I is random, supported on 1,...,10 days, with probability distribution $\Gamma(5, 1)$

[0.0037, 0.0497, 0.1341, 0.1893, 0.1912, 0.1578, 0.1138, 0.0745, 0.0454, 0.0261]

3. "infectiousness" $\beta(n)$ on day $I + n$ is given by (the same vector for now!)

[0.0037, 0.0497, 0.1341, 0.1893, 0.1912, 0.1578, 0.1138, 0.0745, 0.0454, 0.0261]

Therefore expected infectiousness $e(n)$ on day n after infection is "proportional" to $\Gamma(10, 1)$. We approximate this by the following discrete distribution (supported on $[1, 20]$):

[0.0000, 0.00106, 0.00705, 0.0238, 0.0522, 0.0858, 0.114, 0.130, 0.130, 0.118,
0.0984, 0.0768, 0.0566, 0.0397, 0.0266, 0.0172, 0.0108, 0.00654, 0.00387, 0.00224]

4. Over our 14-day period, (indexed starting on Monday of an "on" week as day 1), on day d the rate $r(d)$ at which our worker gets infected is $r(d) = c \cdot p(d)$, which varies as follows, in the base case, repeating periodically with period 14:

$$[c \cdot p_{work}, c \cdot p_{home}, c \cdot p_{home}, \\ c \cdot p_{work}, c \cdot p_{home}, c \cdot p_{home}]$$

Here c is a constant depending on the prevailing rate of infection in the population, and the constants p_{work} and p_{home} are just supposed to be a function of the types and amount of contact that our worker has with potentially infected others on a work day or on a day at home, respectively. The constant c will come out in the wash when we compare different rotas.

5. In the modified one week on / one week off rota, $r(d)$ follows the repeating pattern

$$[c \cdot p_{work}, c \cdot p_{home}, c \cdot p_{home}, \\ c \cdot p_{home}, c \cdot p_{home}].$$

6. To investigate the rotating high contact / low contact scenario, in the base case our worker's daily infection risk $r(d)$ might follow the repeating pattern

$$[c \cdot \frac{p_{high} + p_{low}}{2}, c \cdot p_{home}, c \cdot p_{home}, \\ c \cdot \frac{p_{high} + p_{low}}{2}, c \cdot p_{home}, c \cdot p_{home}].$$

while in the modified schedule it follows the pattern

$$[c \cdot p_{high}, c \cdot p_{home}, c \cdot p_{home}, \\ c \cdot p_{low}, c \cdot p_{home}, c \cdot p_{home}].$$

7. If day d is the n^{th} day after infection, the expected infectiousness of our worker on day d is $e(n)p(d)$. Or to put it another way, if our worker is infected on day d then on day $d + n$ her infectiousness is $e(n)p(d + n)$.
8. Notice we are assuming a symmetric contact scenario, so that the rate at which the worker gets infected and the rate at which she infects other people vary in the same way, up to a multiplicative scalar, depending on whether she is at home, at work, in the high-contact role, or in the low-contact role.
9. Therefore the expected number of people that our worker eventually directly infects, *per fortnight of exposure*, is

$$\mathbb{E}(\text{others infected}) = c \sum_{d=1}^{14} p(d) \sum_{n=1}^{20} e(n)p(n + d).$$

10. This scales linearly in c , but our aim here is only to compare two different rotas with comparable values of $\sum_{d=1}^{14} p(d)$, say, so the overall scaling constant c can be ignored.

Take for example $p_{work} = 0.6$ and $p_{home} = 0.02$, meaning that both the infection risk and the opportunity to pass the infection on are five times as high at work as they are at

home. We compare the base case (working both weeks Monday to Friday) with the 5 days on / 9 days home rota. In the base case we have

$$\mathbb{E}(\text{others infected}) = 2.6652c.$$

In 5 days on, 9 days off

$$\mathbb{E}(\text{others infected}) = 0.5241c.$$

In the *alternating days* rota, our worker is at work on Monday, Wednesday, Friday on week 1 then Tuesday, Thursday on week 2. (And the worker they job share with does the same but shifted by a week so that they dovetail.) Then we get

$$\mathbb{E}(\text{others infected}) = 0.727c.$$

So this is not as good as 1 week on, 1 week off, but it still substantially less than half of the base case.

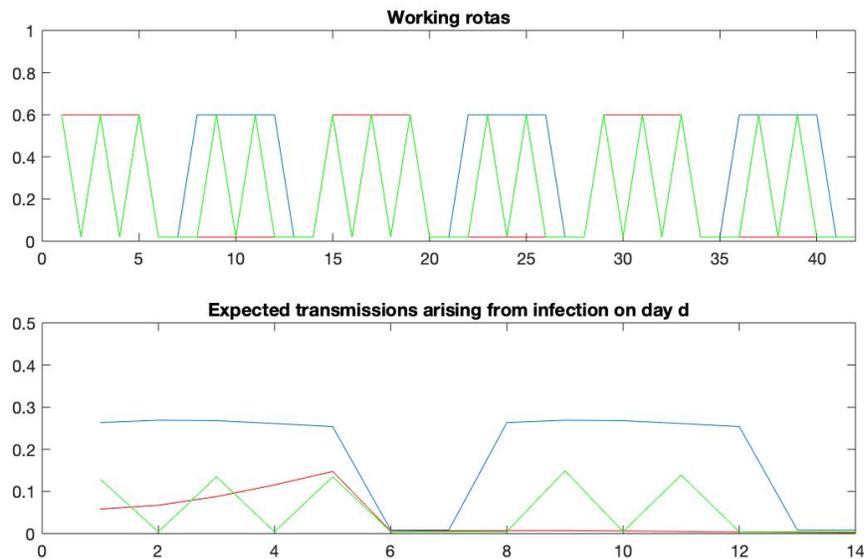


Figure 7: Colours: red is five days on / 9 days off, blue is base case and green is alternating days. For clarity, the plot is not showing expected transmission "conditional" on our worker being infected on day d . It is small for every day that is spent at home because the risk of becoming infected there is very small for the above parameters.

Figure. 7 shows a comparison between these three rotas: the x-axis is indexed by days of the fortnightly repeating period. The y-axis in the top plot shows $p(d)$ for each rota (colours are ...). The y-axis in the lower plot shows the expected number of transmissions resulting from events in which the worker gets infected on day d .

For the rotating roles problem, take for example $p_{high} = 1.0$, $p_{low} = 0.2$, so that $\frac{p_{high}+p_{low}}{2} = 1.0$ and $p_{home} = 0.02$. The "averaged" version where the worker does each role for half of each day is the base case again (so we get $2.6652c$). But with week-high / week-low rotation (with these numbers) we get

$$\mathbb{E}(\text{others infected}) = 2.2764c.$$

This 14% lower than the base case, which has the same value of $\sum_{d=1}^{14} p(d)$. This is shown in Figure 6.

Conclusions and caveats

Our model is very simplistic, and we only considered one-step infections; we did not take account of the fact that a rota system might itself make the daily risk of being infected at work be non-uniform, due to transmission between co-workers. The model is most appropriate when the infection and transmission risk are mainly to members of the public rather than colleagues, in which case we would expect the second order effect of rotas on transmission between coworkers to be relatively small.

It looks as though in the situation where the risk of infection on days off is very low compared with the risk of infection at work, it is substantially better for a job to be shared between part-time two workers who alternate weeks than it is for a single worker to do the job full-time and the other worker be furloughed. The 5 days on / 9 days off rota looks somewhat better than the alternation of working days.

The effects were much less marked when the ratio of risk at work to risk at home was not so high (for example a ratio of 2 or 3).

It is not clear that the sort of flexibility in rota assignment that we have assumed is available to a very large proportion of the workforce; however this could offer an improvement during a time when businesses return to operating half of the time or at half-capacity.

A.2 Low degree interactions / wide isolationsphere (contact tracing)

“Minimising Total Contacts” Principle

Outline conclusion: with rigorous contact tracing and interactions limited to only a few people per person, infections should not spread far. With more than a few interactions, it is unlikely that contact tracing will prevent an infection spreading through an entire cohort.

In a closed workplace such as an office where we can control exactly who interacts with who, and where there is a policy of “if person X gets infected, everybody in their n -neighbourhood is isolated” (see Figure 8), then we can limit the chances of an outbreak infecting an entire cohort by reducing the number d of other people each person interacts with. The more rapidly that employees report having symptoms, the larger d can be, but it must never be allowed to become very large.

Estimates for d for some plausible-sounding parameters are given in the rationale section. It does not get very big ($d = 6$ even in favourable circumstances and with $n = 2$ levels of isolation), suggesting this approach might not be useful in schools, which have high-degree interactions. It grows, but not very quickly, with n .

We considered but did not investigate including strategic testing in our model.

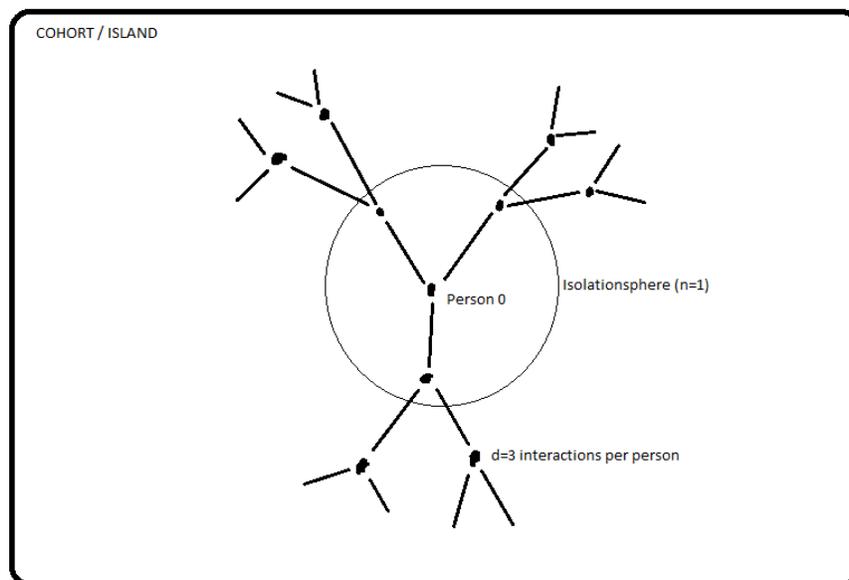


Figure 8: Dots are people, edges are allowed interactions. We assume that when person 0 discovers they are ill that everybody in the circle is isolated, and therefore cannot infect anybody else. We aim to control the number of infections that occur just outside the circle.

Caveats:

- This should be used in combination with cohorting (i.e. disconnecting the graph), since it only limits the chances of an uncontrollable outbreak.
- We took a simplistic model that assumed each person-to-person interaction has an equal rate of infection, and that the rate is gamma distributed. Actual realistic parameters are not known.
- We assumed the original infected individual is aware they are infected, and reports it, after a fixed amount of time.

Rationale:

- At time 0, person 0 gets infected.
- They interact with d people, who each interact with d people, etc... This makes a connection graph of degree d .
- They realize they are infected after R days (R for realize). R is a random variable, but we'll treat it as fixed for now (e.g. $R = 6$).
- At this point, everybody n steps away in the graph is isolated, and we assume therefore cannot infect anybody else. If $n = 0$ it means only person 0 is isolated. If $n = 1$, person 0 and their contacts are isolated. This is the "isolation sphere".
- When person i is infected, and they interact with person j , they become infected after $T_{i,j}$ days, unless the person i was isolated first.
- Let's assume $T_{i,j}$ are i.i.d. (reasonable, but does not account for different strengths of interaction)
- Let's assume they are $\Gamma(k, \theta)$ distributed.
- Then a person $n+1$ steps away from person 0 (i.e. just outside the "isolation sphere") gets infected after time $\Gamma((n+1)k, \theta)$ distributed (unless an isolation prevented it).
- The probability p_n one such person is infected is the probability a sample of this distribution is less than R , so $p_n = \text{cdf}(\Gamma((n+1)k, \theta), R)$.
- There are d^{n+1} such people.
- The expected number of infections just outside the "isolation sphere" is $p_n d^{n+1}$, so for the outbreak to not spread uncontrollably we need this less than 1. (Note: there is independence between infections, how does that change things?)

Mathematical Principles for Unlocking the Workforce - Working Paper

- So require $d < d_n = p_n^{-1/(n+1)}$.
- We also expect $1 + p_0d + p_1d^2 + \dots + p_{n-1}d^n$ infections within the "isolation sphere", and want to keep this small too, presumably.
- For example, if $T \sim \text{Exp}(6) = \Gamma(1, 6)$ and $R = 6$ then $d_0 = 1.6, d_1 = 1.9, d_2 = 2.3$.
- If instead $R = 3$ then $d_0 = 2.6, d_1 = 3.3, d_2 = 4.1$.
- If $T \sim \Gamma(3, 2)$ (same mean as above but more concentrated around the mean, so a longer time before becoming infectious) and $R = 6$ then $d_0 = 1.7, d_1 = 3.5, d_2 = 6.4$.
- If $T \sim \text{Exp}(2)$ then $d_0 = 1.05, d_1 = 1.1, d_2 = 1.2$.
- If $T \sim \text{Exp}(10)$ then $d_0 = 2.2, d_1 = 2.8, d_2 = 3.5$.
- So the larger the "isolation sphere", and the earlier isolation occurs, the larger d can be.
- If $T \sim \text{Exp}(6)$ and $R = 6$ if we were to allow different values for d then n would have to increase accordingly: $(d = 2, n = 2); (d = 3, n = 4); (d = 4, n = 7); (d = 5, n = 10)$.

d	2	3	4	5
n	2	4	7	10
people isolated	7	121	21,845	12,207,031

if $T \sim \text{Exp}(10)$ and $R = 6$ we would rather have: $(d = 2, n = 0); (d = 3, n = 2); (d = 4, n = 3); (d = 5, n = 5)$.

d	2	3	4	5
n	0	2	3	5
people isolated	1	13	85	3,906

if $T \sim \text{Exp}(2)$ and $R = 6$ then we have: $(d = 2, n = 10); (d = 3, n = 18); (d = 4, n = 26); (d = 5, n = 34)$.

d	2	3	4	5
n	10	18	26	34
people isolated	2,047	581,130,733	$6 \times 10^{15} *$	$7 \times 10^{23} *$

* i.e. everyone on the planet

$$\text{people isolated} = \frac{d^{n+1} - 1}{d - 1}$$

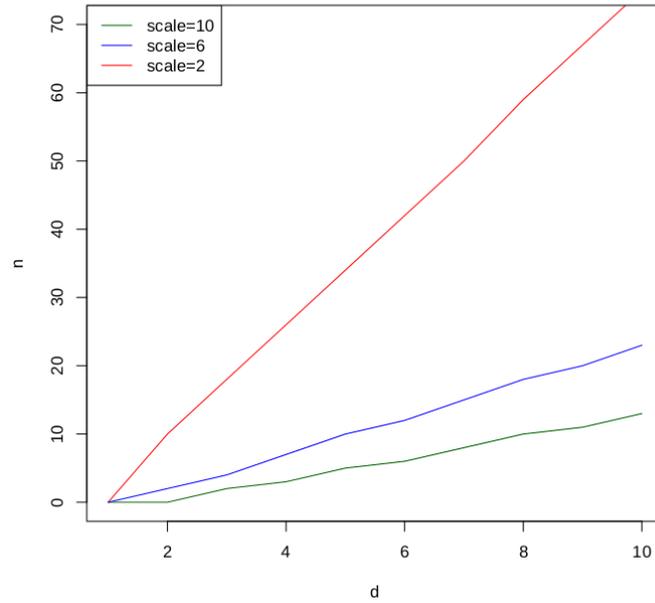


Figure 9: Isolation sphere size versus interaction number for distinct incubation times.

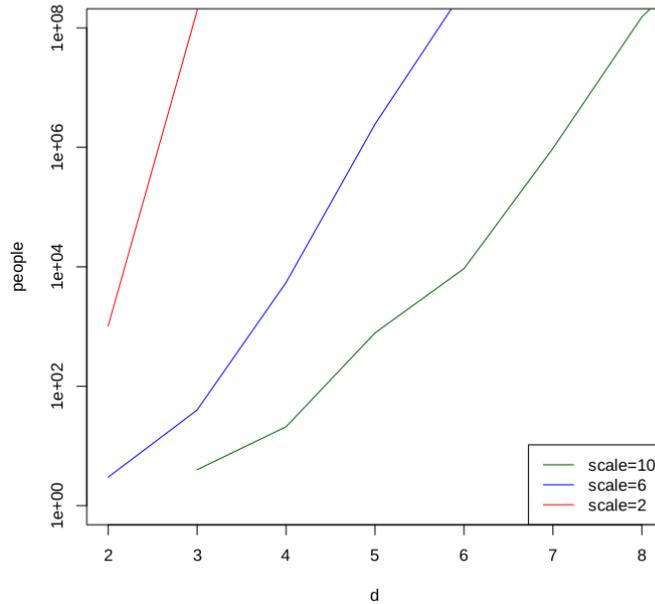


Figure 10: Number of people to be isolated versus interaction number for distinct incubation times.

Simplified model

If one person infects $Po(\lambda)$ people then n th neighbourhood is $Po(\lambda^n)$. So threshold for escaping (comparison of mean to 1) is same in each case, but small case behaviour could differ.

Probability of escaping n degrees of quarantine (everybody within distance n of an initial individual) is approximately $1 - (pd)^{n+1}$. Even if R drops to 0.7 then we would need eight degrees of quarantine to contain an outbreak with 95% probability. At $R = 0.5$ we would need four degrees.

This ignores time delays in transmission. But it tells you that even if you have an environment in which transmission is well controlled, if contact tracing is not performed quickly then extremely aggressive quarantining is necessary to prevent an infection escaping to the general population.

B List of Acronyms

PPE Personal Protective Equipment

V-KEMS Virtual Forum for Knowledge Exchange in the Mathematical Sciences

VSG Virtual Study Group