The benefits of peer transparency in safe workplace operation post pandemic lockdown

Wey, Arkady; Champneys, Alan; Dyson, Rosemary; Alwan, Nisreen; Barker, Mary

DOI:
[10.1098/rsif.2020.0617](https://doi.org/10.1098/rsif.2020.0617)

License:
Creative Commons: Attribution (CC BY)

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (Harvard):
Wey, A, Champneys, A, Dyson, R, Alwan, N & Barker, M 2021, 'The benefits of peer transparency in safe workplace operation post pandemic lockdown', *Journal of The Royal Society Interface*, vol. 18, no. 174, 0617, pp. 1-14. <https://doi.org/10.1098/rsif.2020.0617>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Research

Cite this article: Wey A, Champneys A, Dyson RJ, Alwan NA, Barker M. 2021 The benefits of peer transparency in safe workplace operation post pandemic lockdown. *J. R. Soc. Interface* **18**: 20200617.
<https://doi.org/10.1098/rsif.2020.0617>

Received: 31 July 2020
Accepted: 6 January 2021

Subject Category:

Life Sciences—Mathematics interface

Subject Areas:

biomathematics, computational biology

Keywords:

COVID-19, workplace, mutualism

Author for correspondence:

Alan Champneys
e-mail: a.r.champneys@bristol.ac.uk

The benefits of peer transparency in safe workplace operation post pandemic lockdown

Arkady Wey¹, Alan Champneys², Rosemary J. Dyson³, Nisreen A. Alwan^{4,5,7} and Mary Barker^{5,6}

¹Industrially Focused Mathematical Modelling (InFoMM) EPSRC Centre for Doctoral Training, Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

²Department of Engineering Mathematics, University of Bristol, Bristol BS8 1TR, UK

³School of Mathematics, University of Birmingham, Birmingham B15 2TT, UK

⁴School of Primary Care, Population Sciences and Medical Education, University of Southampton, Southampton SO16 6YD, UK

⁵NIHR Southampton Biomedical Research Centre, and ⁶MRC Lifecourse Epidemiology Unit, University of Southampton and University Hospital Southampton NHS Foundation Trust, Southampton SO16 6YD, UK

⁷NIHR Applied Research Collaboration (ARC) Wessex, Southampton, UK

AW, 0000-0002-3088-9213; AC, 0000-0001-7772-3686; RJ, 0000-0001-6834-5724; MB, 0000-0003-2976-0217

The benefits of different levels of engagement with test, trace and isolate procedures are investigated for a pandemic in which there is little population immunity, in terms of productivity and public health. Simple mathematical modelling is used in the context of a single, relatively closed workplace such as a factory or back-office where, in normal operation, each worker has lengthy interactions with a fixed set of colleagues. A discrete-time SEIR model on a fixed interaction graph is simulated with parameters that are motivated by the recent COVID-19 pandemic in the UK during a post-peak phase, including a small risk of viral infection from outside the working environment. Two kinds of worker are assumed, *transparent*s who regularly test, share their results with colleagues and isolate as soon as a contact tests positive for the disease, and *opaque*s who do none of these. Moreover, the simulations are constructed as a ‘playable model’ in which the transparency level, disease parameters and mean interaction degree can be varied by the user. The model is also analysed in the continuum limit. All simulations point to the double benefit of transparency in both maximizing productivity and minimizing overall infection rates. Based on these findings, public policy implications are discussed for how to incentivise this mutually beneficial behaviour in different kinds of workplace, and simple recommendations are made.

1. Introduction

This study is inspired by the situation in the UK in the latter half of 2020 as the nation has been attempting to restart the economy in the aftermath of the first COVID-19 virus infection peak. The general structure of our mathematical model and the parameter values chosen are specific to that case study. The results are nevertheless intended to be applicable to more general situations in any modern society where there is a residual risk of infection from a virus or other pathogen with insufficient natural immunity in the general population.

There has been much recent evidence to suggest that the most effective containment measure in a human epidemic with relatively small proportions of infectious individuals is that of rapid testing, contact tracing and isolation of those in the contact group [1,2]. The effectiveness of such a strategy is thought

to be a function of the basic reproduction number of the infection, known as the R -value; this gives, on average, the number of new infections that each infection generates. However, R itself is a crude measure, as its instantaneous value will be a function not just of the basic disease dynamics, but of the behaviour of infectious individuals.

The seminal paper of Keeling & Eames [3] introduced ideas from graph theory to epidemiology, where the nature of interactions between infected and susceptible individuals defines a dynamic contact network. Ideas from modern network science, such as the degree distribution, can then be used to estimate statistical properties of the infection, such as the R -number, and to evaluate the effectiveness of different potential treatment strategies; see [4] for a relatively recent review of the state of the art.

The majority of studies that have looked at contact tracing as an effective means of viral control have considered the question at a general population level. We note the recently published studies [1,2], which model the requirements of an effective testing, tracing and isolation strategy to avoid a second-wave of the COVID-19. One should not, however, underestimate the required effort. For example, based on data obtained in a unique collaboration with BBC, Kucharski *et al* conclude ‘in a scenario where there were 1000 new symptomatic cases that met the definition to trigger contact tracing per day [...] 15 000–40 000 contacts would be newly quarantined each day. [...] A high proportion of cases would need to self-isolate and a high proportion of their contacts to be successfully traced to ensure an effective reproduction number that is below one in the absence of other measures.’ This finding is even more stark when combined with results, e.g. [5], that suggest that social isolation needs to happen sufficiently quickly to be effective.

The most important features of any public health campaign built around testing, contact tracing and isolation is the degree of compliance in the general population, which can vary with the method used [6]. Evidence presented to Scientific Pandemic Influenza group on Behaviour (SPI-B) as part of the advice offered to the UK Government Scientific Advisory Group for Emergencies (SAGE) indicated a shortlist of factors that might help to promote compliance with and, adherence to, all behaviours that minimize transmission of SARS-CoV-2 infection. These apply equally to compliance with testing and contact tracing. Factors included messaging that increased perceptions of risk, clear communications from Government identifying what behaviours the public should adopt, encouraging support from the community so creating social norms for infection-limiting behaviours and importantly, actions making it as easy as possible for people testing positive to isolate [7–9]. Recent studies have indicated, however, that other factors may undermine compliance with contact tracing. Perceived lack of data security and privacy, together with lack of trust in government, were found to be the main barriers to adoption [10]. Even in relatively compliant populations, contact tracing may not be sufficient to control the spread of the virus. For example, the study [11] considered the effectiveness of two different methods of contact tracing within a closed, and generally compliant population, namely the participants at a scientific conference on epidemic modelling. One approach was based on reported contacts, as recorded in a log book, the other based on the use of unobtrusive wearable proximity sensors. While both methods were highly tolerated, it was found that neither on

its own was able to give a full picture of meaningful contacts that might have caused an infection to spread. It is clear that other methods are needed to control the spread of the virus in addition to contact tracing.

The question addressed in this paper is more modest. We imagine a risk assessment is to be made within a specific workplace on whether and how the workplace can be made ‘safe’ to reopen following lockdown. Here we use the term *safe* to refer to the public health of the whole of society, to avoid the workplace contributing to a resurgence of the virus in the general population. Nevertheless, the employees, who are the agents in our model, also benefit from safety, but it is assumed that the individual mortality and morbidity rates are sufficiently low that the pay-off to the individual is small. In addition, stakeholders who benefit from the output of workplace will want the workplace to be productive as possible. Such stakeholders include general actors who benefit from the upturn in economy, the owners or shareholders of the business in question, and the workers themselves in terms of security of employment.

Thus, it might seem that safety and productivity are potentially conflicting aims. Given a small overall rate of viral infection, an employer might seek to maximize the workplace productivity by staying open, without isolation of exposed workers. Such actions would clearly compromise safety, and are thus negative to society as a whole. This might be couched in terms of the classical Prisoner’s Dilemma problem within game theory. That is, if every workplace took this attitude, then clearly resurgence within the general population would be probable, whereas one or two isolated ‘bad apple’ employers might be able to benefit by maximizing their personal productivity, provided that others do not.

In fact this ‘bad apple’ principle has been analysed in the context of epidemics by Enright & Kao [12], who ran an agent-based simulation where there are precisely such conflicting pay-offs. The specific motivation for their study was disease among farm animals where the disincentive to the farmer of complying with safety might be particularly harsh, such as the slaughter of their entire herd. They found that a sharp phase transition occurs from sub-critical R values to super-critical, for a relatively small amount of compliance. Their findings are echoed by those of Eksin *et al* [13] in a stochastic network simulation of epidemics. The latter also conclude that ‘a little empathy goes a long way’, meaning that a focus on treating and isolating those infected can be more advantageous than an approach that seeks to protect the healthy. See also the recent review [14] on the state of the art for game-theoretic approaches to analysing agent behaviour within epidemics.

The present study was motivated by discussions during a virtual study group workshop on Mathematical Principles for Unlocking the Workplace at the end of April 2020, organized by the VKEMS¹ initiative between the Isaac Newton Institute (Cambridge), the International Centre for Mathematical Sciences (Edinburgh) and the UK’s Knowledge Transfer Network. The participants were split into four teams to discuss different natures of the person-to-person interactions that happen in different working environments; brief versus lengthy, and with a closed set of colleagues versus with an open set of clients. This study arose from the group look at working environments where interactions were lengthy but with a closed set of colleagues.

The kind of workplaces we have in mind then, are those where workers have a relatively static interaction network

Figure 1. Sketch outlining the operation of the discrete-event model. Here, S is susceptible, E is exposed (not yet infectious), A is asymptomatic infectious, U is unwell infectious, Q is quarantined and R is recovered (and immune). Lines/arrows give all possible transitions between states. See text for details.

with other colleagues. Any interactions with other employees or clients can be assumed to be socially distanced. An example of such a working environment might be a back office, which is split into teams that are physically co-located with a number of middle managers and service personnel who naturally migrate between several teams. Another example might be a factory where individual job functions are well demarcated and a typical worker would only need to interact with a relatively small subset of other workers as part of their normal duties. Thus we shall make the simplifying assumption that the workplace can be represented by a fixed interaction network, with each worker as a node, connected by links that represent interactions between workers who by the nature of their job function cannot effectively socially distance from each other.

A recently published study [15] performs agent-based simulations of the spread of the COVID-19 epidemic within a closed ‘facility’ rather than an open population. We note though that this work does not address the specific question addressed here, namely the effect of having some proportion of the population self-isolate, that is remove themselves from the facility, in the case that they have been in contact with someone who develops symptoms. Also, we do not consider a completely closed facility. Each of our workers are assumed to go about their daily business outside of the workplace with some base level of infection rate each day.

In the present work, we presume a world in which rapid testing, with near instantaneous results, is available to anyone displaying symptoms. Moreover, we assume a routine contact tracing system is in place which requires contacts of those testing positive to self-isolate. We suppose that there is nonetheless a risk of infection outside of the workplace. We wish to explore the question of the effect within our chosen workplace of measures designed to stop any incoming infection spreading to the entire workforce. In particular, we shall look at the effect of the presence of a proportion of opaque workers, who are not transparent about their infection-risk status, and do not go home when made aware that one of their colleagues has the virus. Such opacity may be interpreted as a form of presenteeism, where workers continue to go to work despite possible infection or symptoms. Is there an incentive for the employers and employees alike to engage fully in transparency? That is, is it valuable to engage with

test, trace and isolate procedures in order to halt virus spread while also maximizing productivity?

The rest of the paper is outlined as follows. The subsequent section introduces our mathematical model, its underlying assumptions and the various parameters that may be tuned to simulate different scenarios. Section 3 contains simulation results under the two scenarios where the underlying rate of infection in the general population is either negligible or significant. We also conduct some approximate mathematical analysis to help explain the results. Section 4 contains discussion of the findings both from a scientific perspective and from the point of view of public policy interventions and workplace psychology. Finally, §5 makes recommendations.

2. The mathematical model

We have developed a simple discrete-time simulation model posed on a graph representation of an office environment. An outline of the model is given in figure 1.

2.1. Underlying assumptions

The basic disease model we choose is a form of SEIR model. That is, the usual extension to the Kermack–McKendrick model that allows for four states; (S)usceptible, (E)xposed (infected but not yet infectious), (I)nfectious and (R)ecovered (with immunity from re-infection). In line with what is known about COVID-19, in fact, we choose a modified variant in which there are two infectious states; A for (A)symptomatic infectious individuals and U for infectious symptomatic, or (U)nwell. We suppose that susceptible individuals can become exposed in one of two ways, either through a (small) probability of exposure to the virus present in the general population, or with a much larger probability if one or more of their connected co-workers is infectious. For simplicity, we shall suppose that the rate of exposure outside of the workspace is constant, irrespective of each worker’s personal circumstances. For the purposes of the simulation, it will also be useful to consider a sixth disease state, (Q)uarantined, to represent those who are obliged to be in quarantine post infection. Note that the Q state is only a small subset of those who are isolating at home (figure 1), which also includes other disease states.

Thus we can represent the disease dynamics as

$$S \rightarrow E \rightarrow A \rightarrow U \rightarrow Q \rightarrow R,$$

with the possibility for remaining in each state for periods of time and the possibility of skipping certain stages altogether (for example, A may transition straight to Q or to R without going through U). The full set of possible transitions are illustrated by the arrows in figure 1. The specific disease parameters and logical rules that determine the transitions between these disease states are presented in §2.2.

For convenience, we choose a discrete-time version of the model, in which the fundamental unit of time is the working day. To correct the model for the effect of weekends or other regular workplace closure days, we could in principle exclude such days from our simulation and choose α to be a given function of time, which would be larger after a closure day, and adjust the time intervals $t_{E,A,U,Q}$ accordingly.

The fundamental model is a dynamic network with N nodes, in which nodes are workers and the state at node i is a 3-tuple:

$$\begin{aligned} X_i &= (x_i, p_i, o_i); \\ x_i &\in \{S, E, A, U, Q, R\}, \\ p_i &\in \{0, 1\}, \\ o_i &\in \{0, 1\}. \end{aligned}$$

Hence, x_i gives the disease state; p_i is a binary variable that measures whether the worker is present in the workplace ($p_i = 1$) or is self-isolating at home ($p_i = 0$); o_i is a binary variable that determines the *opacity* of the worker, namely whether they consistently engage in test, trace and isolate and share their data openly ($o_i = 0$), or not ($o_i = 1$). The workplace contact network of interactions is given by an adjacency matrix

$$A = \{a_{ij}\}; \quad a_{ij} \in \{0, 1\},$$

such that i, j are in contact iff $p_i p_j a_{ij} = 1$.

Note therefore that while we assume A_{ij} is fixed in time, the actual workplace contact network varies according to whether workers are at home or not. It is useful therefore to define a current workplace contact matrix C and the set of *workplace contacts* W_i for each node i :

$$C = \{c_{ij}\} = \{p_i p_j a_{ij}\} \quad W_i = \{j : C_{ij} = 1\}.$$

It is also helpful to define indicator functions f_i and g_i to determine respectively whether one of node i 's contacts is infectious, or whether one is reporting symptoms

$$f_i = \begin{cases} 1 & \text{if } \exists j \in W_i \text{ such that } x_j \in \{A, U\}, \\ 0 & \text{otherwise} \end{cases}$$

and

$$g_i = \begin{cases} 1 & \text{if } \exists j \in W_i \text{ such that } o_j = 0 \text{ and } x_j = U, \\ 0 & \text{otherwise.} \end{cases}$$

The opacity variable o_i determines an individual worker's behaviour if they become symptomatic or if one of their contacts tests positive for the disease.

2.2. Time increment

The key to the model is the time update step, which determines how each worker transitions between disease states. We first delineate all the update rules for transparent

workers, before explaining only what is different in the case of an opaque worker.

The behaviour of transparent workers: assume $o_i = 0$, and consider the possible transitions taking place in one time step to worker i .

- (i) **Infection:** If worker i is in work and susceptible, such that $p_i = 1$ and $x_i = S$, then the probability of infection (transition to disease state E) in the next time step is $\alpha + \beta f_i$. Otherwise x_i remains as S .
- (ii) **Isolation because of contact:** In addition, if worker i is in work, and a transparent work contact is unwell, such that $p_i = 1$ and $g_i = 1$, then $p_i \rightarrow 0$ in the next time step. That is, worker i goes home, irrespective of their own disease state x_i in the present time instant or the next.
- (iii) **Exposure:** If worker i is exposed, such that $x_i = E$, then x_i remains in this state for a total number of contiguous days equal to t_E , before transitioning to disease state A , as i becomes infectious but asymptomatic.
- (iv) **Asymptomatic infectiousness:** If i is asymptomatic, such that $x_i = A$, then x_i remains in this state for a total number of contiguous days t_A . After this, x_i transitions, in the next time step and with probability γ , to U . Otherwise, x_i remains in disease state A for a further t_U days, without ever passing to U . In other words, there is a probability that worker i never feels unwell.
- (v) **Symptomatic infectiousness:** If worker i is unwell, such that $x_i = U$, then i remains in this state for a total of t_U contiguous days, after which worker i passes to the recovery step.
- (vi) **Isolation due to symptoms:** In addition, if worker i is in work and unwell, such that $p_i = 1$ and $x_i = U$, then $p_i \rightarrow 0$ for the current time step. That is, as soon as an individual becomes unwell, they do not come into work (provided that they are transparent).
- (vii) **Recovery:** After $t_A + t_U$ time steps since the transition to disease state A , worker i becomes disease free. If $p_i = 0$, then x_i transitions to Q as i goes into a post-symptomatic quarantine state, in which it remains for a total of t_Q consecutive time steps, after which i goes to the immunity step. Else, that is if $p_i = 1$, worker i passes to the immunity step.
- (viii) **Immunity:** At the end of the infection, with probability δ , the worker develops viral antibodies so that $x_i \rightarrow R$ and $p_i \rightarrow 1$. Else, with probability $1 - \delta$, the worker does not become immune, so that $x_i \rightarrow S$ and $p_i \rightarrow 1$. If $x_i = R$ at any time step, irrespective of the value of p_i , then worker i remains in R for the rest of the simulation, and p_i remains 1.

The behaviour of opaque workers. Workers who are not-transparent, $o_i = 1$, are assigned at the beginning of the simulation and remain that way for all time steps. Their behaviour is identical to that of transparent workers except

- (ii') If $p_i = 1$ and $g_i = 1$, then $p_i \rightarrow 0$ in the next time step, with probability ϵ . Otherwise p_i remains 1. That is, even if present opaque workers have unwell contacts, they might stay in work.
- (iv') If $p_i = 1$ and $x_i = U$, and $x_i = A$ in the previous time step, then $p_i \rightarrow 0$ for the current time step with probability ζ . Otherwise p_i remains 1 for the entire time t_U . So when

Table 1. Parameter values used in the simulations. See text for interpretation and justification.

parameter	meaning	value
N	number of workers	100
T	period of simulation (days)	100
d	mean number of workplace contacts	3–12
O	percentage of nodes that are opaque	0–100
α	probability of community infection	{0, 0.001}
β	probability of infection from an infectious connection	0.1
γ	probability of becoming symptomatic	0.95
δ	probability of gaining immunity after infection	0.5
ϵ	probability of non-transparent isolating due to contact	0.01
ζ	probability of non-transparent isolating due to symptoms	0.01
t_E	incubation period before infectious (days)	4
t_A	initial asymptomatic period while infectious (days)	3
t_U	time following t_A until disease free (days)	7
t_Q	required time of quarantine after symptoms stop (days)	5
μ_1	productivity of home working	{0, 0.7, 1}
μ_2	productivity at work while sick	{1, 0.2, 1}

opaque workers become unwell, they may not go home and thus their contacts are unaware that they may be infected and also do not go home.

2.3. Justification of model choices

The parameter choices in table 1 are chosen to be indicative and consistent with, rather than specifically fit to, current estimates based on evidence from COVID-19. The reason for not applying more precise fitting is twofold. First, we are not attempting to model a particular workplace, region or nation, nor specific details of governmental test, trace and isolate policies. Second, we are trying to elucidate general principles that might equally apply in future pandemics, not to make specific data-driven conclusions for the case of COVID-19.

The particular parameter choices mean that the length of an infection, following incubation, is $t_A + t_U = 10$ working days, which is consistent with the latest evidence for COVID-19 [16]. Note that, due to reported low death rates among otherwise healthy working age populations, we have simplified by assuming that all workers are eventually sufficiently healthy to return to work at the end of the infection; this may be assumed to be a ‘best case scenario’. We have further simplified by assuming that incubation and infection times $t_{E,A,U}$ are deterministic, whereas a more representative simulation would allow these parameters to be chosen from a distribution. We have also supposed that no individual is vaccinated.

At the time of writing, it is not clear what proportion of individuals obtain immunity having had the disease, so we make the reasonable assumption that $\delta = 50\%$ of infectious individuals develop immunity, which lasts for the rest of the simulation run time. The incubation period before infectiousness, the degree to which individuals are infectious before they develop symptoms, and the proportion of individuals that are infectious but never develop symptoms have also not been clearly established. Thus the relevant parameter choices $t_E = 4$, $t_A = 3$ and $\gamma = 0.95$ are intended to be illustrative of what might be the case. Note that the benefit of transparency is particularly sensitive to the choice of γ , the probability of developing symptoms. Although this value is at the lower end of the current estimate of 4%–41% of COVID-19 positive patients being asymptomatic [17], the chosen values of γ can equally be interpreted within the model as there being a probability of 0.95 of an infectious patient being detected within three working days through a combined regime of regularly testing and reporting of suspicious symptoms.

One simplification in the model is that the probability of infection β is assumed to be constant, independent of the number of infectious neighbours, provided at least one neighbour is infectious. This parsimonious assumption, which corresponds to the classical Greenwood model, e.g. [18], represents what is suspected about COVID-19, namely that people in a closed space with inadequate ventilation, in the presence of at least one infective, are more or less equally at risk of infection. Alternative hypotheses could be a binomial Reed-Frost-type assumption (e.g. [19]). To test the sensitivity of the results to such assumptions, we have also implemented the other extreme where β rises linearly with the number of contacts (results not shown). We found that in both the hot and cold runs the results were qualitatively the same as those with fixed probability β .

Another simplification in the model is that we have not implemented specific network topologies that are built around partial ‘bubbles’ of close contacts that interact with other bubbles as infrequently as possible. There is an increasing body of evidence in epidemic modelling for COVID-19 that suggests that the benefit of isolation bubbles significantly decreases when there is a small amount of leakage between bubbles (e.g. [20]). We tested this in preliminary simulations with synthetic small-world-type networks comprised of local bubbles with one or two long-range connections to other bubbles and found that the transmission rate to be remarkably similar to a random network with the same average degree. So, as a parsimonious approximation, we use an Erdős–Renyi graph with a fixed probability of there being a link between any two nodes.

Throughout each simulation run of the model, all parameters are taken as fixed. Clearly, more accurate simulations would need to test sensitivity to parameters changing over time, including the possibility of stochastic parameter variation. Parameter sensitivity is also key. In what follows, we shall consider sensitivity to two key parameters, which represent the primary ingredients under investigation, the degree of opacity O and the average network degree d . We will also conduct Monte Carlo simulation where the specific network topology varies for the each value of d .

Finally, we have not sought to understand dynamical consequences of testing, other than to assume that an accurate

Figure 2. Examples of randomly generated workplace contact networks for six different random initial conditions. Here, $N = 15$, and one-third of the workforce is opaque (solid square) and two-thirds of the workforce is transparent (transparent circle). One individual is assumed to be infected with the virus and is in state E (the orange node labelled 0) whereas all others are susceptible (green). (a–c) Realizations with average degree $d = 4$; (d–f) $d = 8$.

test is available with near-instant result, as soon as a worker becomes symptomatic. We also assume transparent workers whose contacts have a positive test will simply isolate at home for a requisite period, without receiving a test themselves. We note a recent study has suggested that a programme of widespread polymerase chain reaction (PCR) testing of contacts of identified cases may be a more efficient strategy [21]. In future work, it would be interesting to investigate such a strategy, using the modelling framework proposed here.

3. Simulation results

All simulations were carried out in Python using the parameter values given in table 1.

We have chosen a workplace with $N = 100$ employees. In each run, the topology of this workplace is generated as a random (Erdős–Rényi) symmetric graph with probability d/N that a given edge is present such that $A_{ij} = 1$, independent of other edges. By definition we choose $A_{ii} = 0$, hence the average degree of each graph is actually

$$(N - 1) \frac{d}{N} = 0.99d.$$

For a given opacity, O , we use the same process of ensuring that node i has probability O that $o_i = 1$ independently of the value o_j for any other node j . Figure 2 illustrates examples of graphs that are generated in this manner.

When analysing the results of simulations, it is useful to have a measure of productivity, or the proportion of productive work hours

$$\text{Productivity} = \frac{1}{NT} \sum_{0 < t < T} [\mathbf{i}_{\{x_i \neq U, p_i = 1\}}(x_i) + \mu_1 \mathbf{i}_{\{x_i = U, p_i = 1\}}(x_i) + \mu_2 \mathbf{i}_{\{x_i \neq U, p_i = 0\}}(x_i)], \quad (3.1)$$

where $\mathbf{i}$ represents an indicator function and $\mu_{1,2} \in [0, 1]$. We can then define lost productivity via

$$\text{Productivity Deficit} = 1 - \text{Productivity}. \quad (3.2)$$

The rationale behind the parameters μ_1 and μ_2 is that those in work are assumed to be fully productive if not sick and have

fractional productivity μ_1 if sick, whereas those isolating at home do no work if they are sick and have fractional productivity μ_2 if not. In what follows, we shall take two extreme and one balanced measure of productivity.

$$\begin{aligned} \text{'academic': } \mu_1 = 0, \mu_2 = 1; \quad \text{'factory': } \mu_1 = 1, \mu_2 = 0 \\ \text{and 'office': } \mu_1 = 0.2, \mu_2 = 0.7. \end{aligned} \quad (3.3)$$

Note that in the ‘factory’ case takes the extreme limit that the work is so menial that mere presence of the workforce is sufficient to assure productivity, whereas the ‘academic’ case takes the opposite extreme where working from home is equally effective, but work is impossible when sick.

When running simulations, we shall consider two cases, which we refer to as *running cold* and *running hot* depending on whether the overall rates of infection in society are negligible or not:

$$\text{'running cold': } \alpha = 0; \quad \text{and 'running hot': } \alpha = 0.001. \quad (3.4)$$

3.1. Running cold

In a cold run, the chance of an infection from the outside world is negligible, so that the parameter $\alpha = 0$.

We start simulations at time zero with one exposed individual and all other individuals in state S (we assume that the state of the disease in the general population is that there is a negligible number of individuals who are already immune). Examples of such simulations are shown in figure 3.

The simulations demonstrate the benefits of transparency. The left-hand plots show how an infection that starts with one individual at day 0 quickly dies out for a fully transparent workplace. By contrast, when 50% or 100% of the workers are opaque (middle and right panels) the consequence of that initial infection is still present in the workplace after 100 days. Note how, for fixed values of the other parameters, the size of the contact network between $d = 4$ and $d = 10$ makes little qualitative difference. However, there is a significant quantitative difference when not all workers are transparent. Given 50% opacity, the maximum size of the outbreak is such that with $d = 10$ there are about 25 people who are sick at around day 30, rising to almost 40 people with 100% opacity.

Figure 3. Examples of cold runs for different numbers of workplace contacts d and different workforce opacities O against time. In each plot, the red curve gives the proportion of unwell individuals (in state U) and the blue curve gives the proportion who are at home (with $p_i = 0$). Parameter values are: (a) $d = 4, O = 0\%$; (b) $d = 4, O = 50\%$; (c) $d = 4, O = 100\%$; (d) $d = 10, O = 0\%$; (e) $d = 10, O = 50\%$; (f) $d = 10, O = 100\%$.

Figure 4. (a) Proportion of the workforce that become unwell at least once (red lines) or go home at least once (blue lines) as a function of the average contact network size d . Results are shown for three different opacities $O = 0\%$ (dashed lines), $O = 50\%$ (dot-dashed) and $O = 100\%$ (solid). (b) Mean (solid lines), plus and minus one standard deviation (dashed) of the becoming ill proportions. Different colours represent the three opacities $O = 0\%$ (green lines), $O = 50\%$ (yellow) and $O = 100\%$ (red).

In addition to single runs, it is useful to generate a statistical ensemble. We have done this for a range of d and O values, taking 50 repeats for each parameter value. The results are shown in figures 4 and 5, which show the proportion of workers that become unwell (enter state U) and the proportion that go home (have $p_i = 1$) at least once during the 100 day simulation time.

Observing the results against d (figure 5), we note how, for the 100% opacity case with degree 10, almost all of the workforce appear to catch the disease. Thus a certain amount of herd immunity is established in the population (recall the immunity rate $\delta = 0.5$) and this is what causes the infection rate to decrease towards the end of the simulation. The results for 50% opacity are similar. Smaller contact network sizes however result in smaller infections.

Further conclusions can be drawn from the graphs plotted against opacity in figure 5. Here note that, for the case of the highest degree, there is a sharp increase in the proportion of infected individuals for low opacity. For an intermediate degree, the proportion of infected individuals

appears to vary more linearly with opacity, with the sharp increase, if there is one, occurring later, perhaps around 20% (although note the large standard deviation). For the lowest degree value, the number of infected individuals appears to be low, with the sharpest increase at around 45%.

From these simulations, we can also compute the productivity deficit according to (3.1) and (3.2). The results are presented in figure 6. The results for the 'academic' working environment (green lines) are as expected. In this environment, well individuals are equally productive at home as in the workplace. So productivity is greatly enhanced by high transparency. In the case of factories and mixed offices, especially the factory, the productivity curve is more n-shaped and it may seem that optimal productivity can be gained at 100% opacity. But at what price? In this scenario, for the higher degrees, the majority of the workforce have caught the virus, and it seems clear that in reality the factory will need to be closed down, because it has become a hotbed of infection.

Figure 5. Similar to figure 4 but showing variation with opacity for three different values of average degree; $d = 4$ (dashed lines in *a*, green lines in *b*), $d = 7$ (dot-dashed in *a*, yellow in *b*) and $d = 10$ (solid in *a*, red in *b*).

Figure 6. (*a,c*) Productivity deficit as a function of degree (*a,b*) and opacity (*c,d*) for the three different types of working environment—academic (green lines), office (orange) and factory (red). (*a*) Results are shown at three different opacities 0% (dashed lines), 50% (dot-dashed) and 100% (solid). (*c*) Results are shown at three different degrees $d = 4$ (dashed lines), $d = 7$ (dot-dashed) and $d = 10$. (*b,d*) Mean (solid lines), plus and minus one standard deviation (dashed) for the office scenario. (*b*) Results are shown at three different opacities 0% (green lines), 50% (yellow) and 100% (red). (*d*) Results are shown at three different degrees $d = 4$ (green lines), $d = 7$ (yellow) and $d = 10$ (red).

3.2. Running hot

We also performed exactly the same set of simulations (with fresh randomisations, of course) in the case of running hot (cf. (3.4)). Here, the same initial condition is used as for the cold case, with one randomly chosen exposed individual (state E) at $t = 0$ with all others in state S . The difference now is that, with $N = 100$ and $\alpha = 0.001$, approximately every 10 days a new infected individual is likely to enter the workplace. Data that are exactly analogous to those in the cold case are presented in figures 7–10.

It is worth commenting on the different time-course dynamics in figure 7 than figure 3. In the case of a completely transparent workforce (left-hand panels) note how the infection does not die out for the case of running hot. This is because approximately every 10 days a new infection enters the workplace. Thus the number of sick individuals (about three for $d = 4$ and about seven for $d = 10$) remains constant as does the

number of workers at home (about eight and 20, respectively) throughout the simulation. Note that the number who are sick at any one time is roughly d times the number who would be expected to be sick if there was no contact within the workplace. For either 50% or 100% opacity, in the case of low degree (upper middle and right panels), the running hot case causes the infection to last much longer in the workplace, with a much wider peak than the running cold case. The maximum number of infected individuals also rises, compared to the cold runs, which is even more apparent in the case of 100% opacity.

These observations, based on just isolated simulation runs, are born out in general, in the Monte Carlo parametric runs in figures 8–10 which, apart from the value of α , are run under exactly the same assumptions as figures 3–4 in the cold case. The results here are broadly similar to the cold case, but note the dire consequences of large opacity in terms of the spread of the disease.

Figure 7. Same as figure 3 but in the case of hot runs. Examples of hot runs for different sizes of contact network d and workforce opacity O against time. In each plot, the red curve gives the proportion of unwell individuals (in state U) and the blue curve gives the proportion who are at home (with $p_i = 0$). Parameter values are: (a) $d = 4, O = 0\%$; (b) $d = 4, O = 50\%$; (c) $d = 4, O = 100\%$; (d) $d = 10, O = 0\%$; (e) $d = 10, O = 50\%$; (f) $d = 10, O = 100\%$.

Figure 8. Similar to figure 4 but for running hot. (a) Proportion of workforce that become unwell at least once (red lines) or go home at least once (blue lines) as a function of the average contact network sized. Results are shown for three different opacities $O = 0\%$ (dashed lines), $O = 50\%$ (dot-dashed) and $O = 100\%$ (solid). (b) Mean (solid lines), plus and minus one s.d. (dashed) of the becoming ill proportions. Different colours represent the three opacities $O = 0\%$ (green lines), $O = 50\%$ (yellow) and $O = 100\%$ (red).

3.3. Analysis

A simple analysis can estimate the reproduction number r_0 in the case of different opacities and average network degree, along with the size of the infection within the workplace as a function of parameters by passing to a continuum limit. Taking the limit of a sufficiently long time and sufficient network size $T, N \gg 1$, and assuming that $\epsilon = \zeta = 0$ for simplicity, we can apply a standard mean-field approximation to obtain an equivalent system of ordinary differential equations, e.g. [22, ch. 10]. Specifically, taking the limit of continuous time, $\Delta t \rightarrow dt$, and of a large network of homogeneous mean degree d , applying this method yields the following equivalent SIR system:

$$\dot{S} = \frac{(1-\delta)}{\tau} I - \omega \frac{d}{N} \beta SI - \alpha S, \quad (3.5)$$

$$\dot{I} = \alpha S + \omega \frac{d}{N} \beta SI - \frac{1}{\tau} I \quad (3.6)$$

and
$$\dot{R} = \frac{\delta}{\tau} I, \quad (3.7)$$

where S, I and R now represent the numbers of the total workforce in states $S, \{I, U\}$ and R , respectively, as a function of time (dot represents differentiation with respect to a new continuous time variable). Here,

$$\tau = t_A + t_U \quad \text{and} \quad \omega = O/100\%$$

represent the time of infection and proportion of opaque individuals, respectively, δ is the proportion of people who gain immunity upon recovery and d/N represents the chance that two individuals are connected within the workplace (the average degree of the network divided by the total number of individuals). Note that total population is conserved, that is $S + I + R = N$ is constant.

We first non-dimensionalize the model (3.5)–(3.7), scaling populations with the total population N and time with the recovery timescale τ such that

$$(S, I, R) = N(S^*, I^*, R^*), \quad t = t^*/\tau, \quad (3.8)$$

Figure 9. Similar to figure 8 but showing variation with opacity for three different values of average degree; $d = 4$ (dashed lines in (a), green lines in (b)), $d = 7$ (dot-dashed in (a), yellow in (b)) and $d = 10$ (solid in (a), red in (b)) (cf. figure 5).

Figure 10. Similar to figure 6 but for running hot. (a,c) Productivity deficit as a function of degree (a,b) and opacity (c,d) for the three different types of working environment—academic (green lines), office (orange) and factory (red). (a) Results are shown at three different opacities 0% (dashed lines), 50% (dot-dashed) and 100% (solid). (c) Results are shown at three different degrees $d = 4$ (dashed lines), $d = 7$ (dot-dashed) and $d = 10$. (b,d) Mean (solid lines), plus and minus one standard deviation (dashed) for the office scenario. (b) Results are shown at three different opacities 0% (green lines), 50% (yellow) and 100% (red). (d) Results are shown at three different degrees $d = 4$ (green lines), $d = 7$ (yellow) and $d = 10$ (red).

to find

$$\dot{S}^* = (1 - \delta)I^* - r_0 S^* I^* - a S^*, \quad (3.9)$$

$$\dot{I}^* = a S^* + r_0 S^* I^* - I^* \quad (3.10)$$

and

$$\dot{R}^* = \delta I^*, \quad (3.11)$$

where

$$r_0 = \omega d \beta \tau \quad \text{and} \quad a = \alpha \tau. \quad (3.12)$$

Note that r_0 is the basic reproduction number of the outbreak, representing the average number of new infections one infected person generates within the workplace. The parameter $a = \alpha \tau$ is the balance of timescales between new infections entering the workplace and infected individuals recovering (recall $a = 0$ is the ‘running cold’ scenario such that no new infections are entering the workplace). Along with the immunity rate δ , which is already dimensionless,

these three parameter groupings control the behaviour of the infection within the workplace and will determine the dynamics of the disease. Total population is still conserved, now scaled such that $S^* + I^* + R^* = 1$.

The immediate benefit of non-dimensionalization is that the relative importance of the parameters of the simulation becomes obvious. For example, r_0 depends on the product of d and ω , which implies the opacity that a workplace can tolerate before $r_0 > 1$ is inversely proportional to d the average number of interactions per worker. This duality between O and d is also apparent in the simulation results (cf. the red curves in figure 4 with 5, and figure 8 with 9).

Moreover, steady states of the system (3.9)–(3.11) indicate what we expect to happen for long times. This varies depending on the values of parameters r_0 , a and δ .

When $a > 0$ is non-zero (the running hot scenario) with some level of immunity $\delta \neq 0$, the steady state is $S_0^* = I_0^* = 0$,

Figure 11. (a) Plot of I_0^* given by (3.14) against r_0 for different $a \geq 0$ in the case of no immunity ($\delta = 0$). (b) Positive value of R^∞ obtained by solving (3.15) for different values of $\delta > 0$ in the running cold case ($a = 0$).

$R_0^* = 1$, that is everyone eventually catches and recovers from the disease. When recovery confers no immunity ($\delta = 0$), we find the steady state

$$S_0^* = 1 - I_0^* = \frac{r_0 + 1 + a - \sqrt{(r_0 - 1 - a)^2 + 4ar_0}}{2r_0} \quad (3.13)$$

and

$$I_0^* = \frac{r_0 - 1 - a + \sqrt{(r_0 - 1 - a)^2 + 4ar_0}}{2r_0}. \quad (3.14)$$

In this case, the infection is always present within the workforce with a level dependent on both the infection dynamics on the network (r_0) and the rate of infection outside the workplace a . (Note that this collapses to $I_0^* = 1 - 1/r_0$ when $a = 0$ as would be expected, showing that the infection will die out if $r_0 < 1$.) Graphs of I_0^* versus r_0 are shown in the left-hand panel of figure 11.

For the running cold scenario, such that $a = 0$, with some level of immunity $\delta > 0$, the non-trivial steady state is given by

$$S_0^* = 1 - R^\infty, \quad I_0^* = 0 \quad \text{and} \quad R_0^* = R^\infty,$$

where R^∞ satisfies the transcendental equation

$$r_0(R^\infty - 1) - \delta + 1 = (1 - r_0 - \delta)e^{-r_0 R^\infty / \delta}. \quad (3.15)$$

Note that $R^\infty = 0$ is always a solution to (3.15) for all $r_0, \delta > 0$, but there are also non-trivial solutions, which we plot in figure 11 as a function of r_0 , for various $0 < \delta < 1$. There is a (transcritical) bifurcation at $r_0 = 1$ such that for $r_0 < 1$ the non-zero steady-state value of R_0^* is $R_0^* = 0$, implying that the infection dies out. For $r_0 > 1$, the non-trivial value of $R^\infty > 0$ becomes the stable steady state. Note how the R^∞ curve rises steeply with r_0 , thus explaining the shape of infection curves in figures 4 and 5.

4. Findings from the model results

Although we have only run the simulations for a fixed set of parameters, the analytical results suggest a certain universality to our findings. In particular, there is an inverse proportionality between opacity and average contact degree (since $r_0 = \omega d \beta \tau$); the greater the average connectivity within the workplace, the smaller the opacity must be to avoid the infection taking hold and eventually reaching the potential non-trivial equilibria described in §3.3.

In order to maximize productivity, figures 6 and 10 also show that, in most types of workplace, it is advantageous to have almost complete transparency. This is despite the fact that this involves sending every worker home as soon as they are in contact with an infectious colleague. The exception is the case that we have called a ‘factory’, which is the extreme case where there is no work done from home and productivity is simply the same as attendance. Here, of course, it is best to send no one home, so that optimal productivity is obtained at an opacity of 100%. However, we have shown that this case leads to the fastest possible spread of infection and the largest infected population within the workplace. In reality, if contact tracing is being conducted throughout the general population, such a scenario is likely to lead to the workplace being identified as a hotbed of infection, leading to the ‘factory’ being forced to shut down, negative publicity for the employer and possible prosecution. Also, recall that the ‘factory’ scenario unrealistically ignores the loss of productivity due to sick workers. Recalling the ‘n’-shaped nature of the productivity deficit versus opacity curve, especially in the case of running cold (figure 6), it would seem that an optimal strategy to maximize productivity, even in the ‘factory’ case, would be again to maximize transparency within the workforce.

It might be useful to reflect on what the variable ‘opacity’ really represents. The opacity of the workforce could be construed as the availability of regular testing to that workforce. We have said that transparent workers go home when they enter state U , that is, they first develop symptoms. Instead, we could alternatively consider the state U to be the return of a positive test. Then t_A represents the length of time between catching the virus and testing positive. In such a scenario, rather than thinking of workers as having either helpful or unhelpful behaviours, we can consider transparency as the degree to which a rapid, accurate, regular testing regime is undertaken in the workplace.

The key finding from our simulations then is that a policy of maximum transparency is optimal not only for stopping disease spread (which is unsurprising) but also in terms of maximizing productivity of the workplace. These results suggest that making a workplace safe to reopen in the post peak phase of a pandemic such as COVID-19 requires the adoption of a number of changes to the running of the workplace and new behaviours on the part of both employers and employees.

For example, employees will presumably have, as a condition of employment, to declare to their employers when

they have developed symptoms of the virus or tested positive. They would also need to be incentivized to identify those in the workplace with whom they have been in contact and to quarantine at home. Employers will have to facilitate these communications and put in place arrangements to back-fill posts left empty through sickness absence and quarantine. In a running cold scenario, where infection in the general population is sparse, there could additionally be a requirement for employees to report when they have been exposed to the virus outside of the workplace; to keep the workplace safe, workplace contacts of this employee could also be required to quarantine.

There are well-defined and evidence-based behavioural science principles that can be used to inform how to support and encourage required changes in behaviour. These have been summarized as they apply to reducing infections during the pandemic in a number of recent reviews and commentaries [23–25] and have obvious application to plans to make return to work safe. They suggest that to be effective, a workplace campaign would need to: create a collective viewpoint emphasizing how people can look after each other; ensure messages concerning changes to behaviour come from trustworthy and credible sources; and ensure that whatever employers and employees have the *capability*, opportunity and *motivation* for what they are being required to do.

It would be interesting in future work to extend the modelling framework in this paper to include the modern idea of mechanism design from economic game theory, see for example [26] for application of these ideas to epidemic modelling. In particular, we could seek to enumerate the benefit of different mechanisms that might incentivize transparency in the workplace. That work, which is beyond the scope of this study, should presumably be accompanied by empirical or ethnographic research.

5. Conclusions, behavioural implications and tentative public policy recommendations

The main model finding is that in all situations there is a benefit of what we have called ‘transparency’ in the safe and productive operation of workplaces during a pandemic, and the relationship we have uncovered between the degree of transparency and the safe size of the average working group.

A key question is how one might ensure that such transparency is enacted in practice. It is well-established within behavioural science that incentives offered to compensate for potential losses may be particularly effective as motivators, [27]. The independent SAGE committee report [28] on how best to support effective application of testing for the virus and contact tracing recommends that support for isolation should be provided. Concerns about social disapproval and fairness are likely to interact with, and have reinforcing effects on, compliance with requirements for transparency and quarantining [29].

It is also well established that compliance is also more likely if employers and employees feel a sense of ownership and control over the way that a workplace scheme is designed and run. A perception that one has control over aspects of the workplace and job role has long been linked to increased sense of well-being, lower perceived stress and

better general health (e.g. [30]). Recent evidence review suggests that key to adherence to quarantine are clear understanding of the disease and quarantine procedures, social norms and perceived benefits of quarantine, perceived risk of the disease and practical problems such as maintaining supplies and financial consequences of being out of work [31]. In light of this, government advisory bodies have issued guidance to encourage quarantining that includes emphasizing civic duty, advertising the changing social norms and allowing others in the community to express disapproval and stressing the value of the organization, in this case, the employing organization.

In the light of these observations and the scientific conclusions in the previous section, it would seem helpful to make some tentative recommendations directed at policy makers, both government and employers, as well as to employees in the kind of closed, fixed-interaction workplaces envisaged in this study.

Organize workplaces into small, intersecting groups. We recommend that all workplaces should seek to minimize the average number of work contacts per worker. Other studies have suggested work should take place in fixed ‘bubbles’ that do not interact with each other. For most operations, strict bubbles are not feasible. Instead, our results show that the necessary degree of transparency (the rigour to which test, trace and isolate is required) is directly proportional to the size of the average workplace interaction network.

The no detriment principle. We recommend workplaces engender a culture where there is no perceived detriment to a worker being transparent. One issue can be a workplace where there are a lot of self-employed or ‘gig economy’ workers who only get paid when they are present. Mechanisms need to be established so that all workers are fully remunerated if they are required to self-isolate and penalized if they are found to be at work while infectious. One such policy intervention could be the introduction of governmental statutory sick pay for all workers irrespective of their contractual status, from day 1 of self-isolation.

The benefit of mutualism. As described above, we need to ensure that workplace campaigns to encourage transparency are effective. A sense of shared ownership in maintaining an infection-free workplace would need to be inculcated through a sense of collective efficacy. This might involve an employee/employer partnership that runs its own workplace test, trace and isolate system. This may be more effective than a scheme administered by an outside authority where there is the potential for lack of trust, for example over data privacy.

Financial and other disincentives for non-compliant workplaces. For some businesses, there may be a temptation for the employer to seek short-term profit over long-term benefit. If infection rates are low in the general population, there may be unscrupulous or ignorant employers who try to keep everyone in work. We saw this in our models for the ‘factory’ scenario, in which theoretical maximal productivity could be achieved with no transparency whatsoever. It should be the function of a health and safety inspectorate to provide large penalties to deter such behaviour which would clearly not be in the public good.

Personalized solutions for individual workspaces The quantitative findings in this study are a result of the generic parameter choices made in the simulation runs of our mathematical model. Some of these are disease parameters, which for COVID-19 remain unclear at the time of writing. Others related to the nature of the work; for example, what exactly constitutes the social distance between two workers can greatly affect the probability of infection between individuals. We recommend that individual employers should be encouraged to run a playable version of the mathematical model we have introduced in order to test the safety and feasibility of their operation. This may also involve ethnographic studies both to observe how interactions happen in reality and also to design the optimal psychological and policy interventions to obtain the desired outcome.

Data accessibility. All code and data to reproduce the figures in the document, as well as to run other cases, is available at <https://bitbucket.org/ArkadyWey/workplace-network-model/src/master/>.

Authors' contributions. A.R.C., R.J.D. and A.W. performed the mathematical modelling. M.B. and N.A.A. helped frame the study and provided all the public health and behavioural sciences expertise. All coding was performed by A.W. A.R.C. wrote the first draft, and all authors contributed equally to final drafting.

Competing interests. We declare we have no competing interests.

Funding. The authors thank EPSRC, for their support of industrial and applied mathematics, in particular through the Industrially Focused Mathematical Modelling (InFoMM) Centre for Doctoral Training (CDT) at University of Oxford.

Acknowledgements. The authors thank Julia Gog, Thilo Gross, Jessica Enright and Mason Porter for helpful discussions. They would also like to thank all the organizers and participants at the VKEMS study group, from where this work originated.

References

- Kucharski AJ, Klepac P, Conlan AJK, Kissler SM, Tang W, Fry H, Gog JR, Edmunds JW. 2020 Effectiveness of isolation, testing, contact tracing and physical distancing on reducing transmission of SARS-CoV-2 in different settings: a mathematical modelling study. *Lancet. Infect. Dis.* **20**, 1151–1160. (doi:10.1016/S1473-3099(20)30457-6)
- Royal Society DELVE Initiative. Test, trace, isolate, 2020. See <https://rs-delve.github.io/reports/2020/05/27/test-trace-isolate.html> (accessed 1 June 2020).
- Keeling MJ, Eames KT. 2005 Networks and epidemic models. *J. R. Soc. Interface* **2**, 295–307. (doi:10.1098/rsif.2005.0051)
- Pastor-Satorras R, Castellano C, Van Mieghem P, Vespignani A. 2015 Epidemic processes in complex networks. *Rev. Mod. Phys.* **87**, 925. (doi:10.1103/RevModPhys.87.925)
- Pereira T, Young LS. 2015 Control of epidemics on complex networks: effectiveness of delayed isolation. *Phys. Rev. E* **92**, 022822. (doi:10.1103/PhysRevE.92.022822)
- Read JM, Edmunds WJ, Riley S, Lessler J, Cummings DA. 2012 Close encounters of the infectious kind: methods to measure social mixing behaviour. *Epidemiol. Infect.* **140**, 2117–21130. (doi:10.1017/S0950268812000842)
- Scientific Advisory Group for Emergencies (SAGE). SPI-B: How behaviour may change following testing—initial response, 22 April 2020. See www.gov.uk/government/publications/spi-b-how-behaviour-may-change-following-testing-initial-response-22-april-2020.
- Scientific Advisory Group for Emergencies (SAGE). SPI-B: Key behavioural issues relevant to test, trace, track and isolate—summary, 6 May 2020. See www.gov.uk/government/publications/spi-b-key-behavioural-issues-relevant-to-test-trace-track-and-isolate-summary-6-may-2020.
- Scientific Advisory Group for Emergencies (SAGE). SPI-B: Symptom-based contact tracing in comparison to test-based approaches—note, 29 April 2020. See www.gov.uk/government/publications/spi-b-symptom-based-contact-tracing-in-comparison-to-test-based-approaches-note-29-april-2020.
- Altmann M, Zillesen R, Blasone H, Gerdon F, Bach R, Kreuter F, Nosenzo D, Toussaint S, Abeler J. 2020 Acceptability of app-based contact tracing for COVID-19: cross-country survey evidence. Preprint *medRxiv* (doi:10.1101/2020.05.05.20091587)
- Smieszek T, Castell S, Barrat A, Cattuto PJ, White C, Krause G. 2016 Contact diaries versus wearable proximity sensors in measuring contact patterns at a conference: method comparison and participants' attitudes. *BMC Infect. Dis.* **16**, 341. (doi:10.1186/s12879-016-1676-y)
- Enright J, Kao RR. 2015 A few bad apples: a model of disease influenced agent behaviour in a heterogeneous contact environment. *PLoS ONE* **10**, e0118127. (doi:10.1371/journal.pone.0118127)
- Eksin C, Shamma J, Weitz J. 2017 Disease dynamics in a stochastic network game: a little empathy goes a long way in averting outbreaks. *Sci. Rep.* **7**, 44122. (doi:10.1038/srep44122)
- Chang SL, Piraveenan M, Pattison P, Prokopenko M. 2020 Game theoretic modelling of infectious disease dynamics and intervention methods: a review. *J. Biol. Dyn.* **14**, 57–89. (doi:10.1080/17513758.2020.1720322)
- Cuevas E. 2020 An agent-based model to evaluate the COVID-19 transmission risks in facilities. *Comput. Biol. Med.* **121**, 103827. (doi:10.1016/j.combiomed.2020.103827)
- Cevik M, Tate M, Lloyd O, Maraolo AE, Schafers J, Ho A. 2020 SARS-CoV-2, SARS-CoV, and MERS-CoV viral load dynamics, duration of viral shedding, and infectiousness: a systematic review and meta-analysis. *Lancet Microbe*. **2**, e13–e22. (doi:10.1016/S2666-5247(20)30172-5)
- Byambasuren O, Cardona M, Bell K, Clark M-L, McLaws J, Glasziou P. 2020 Estimating the extent of asymptomatic COVID-19 and its potential for community transmission: systematic review and meta-analysis. *J. Assoc. Med. Microbiol. Infect. Dis. Can.* **5**, 223–234. (doi:10.3138/jammi-2020-0030)
- Bayley NTJ. 1975 *The mathematical theory of infectious diseases and its applications*, 2nd edn. High Wycombe, UK: Griffin.
- Anh HJ, Hassibi B. 2013 Global dynamics of epidemic spread over complex networks. In *52nd IEEE Conf. on Decision and Control*, pp. 4579–4585, Florence.
- Abrahams D et al. 2020 Unlocking higher education space—what might mathematics tell us? July. Virtual Forum for Knowledge Exchange in the Mathematical Sciences white paper. See <https://www.vkemsuk.org/activities/virtual-study-groups>.
- Grassly NC, Pons-Salort M, Parker EPK, White PJ, Ferguson NM. 2020 Comparison of molecular testing strategies for COVID-19 control: a mathematical modelling study. *Lancet Infect. Dis.* **20**, 1381–1389. (doi:10.1016/S1473-3099(20)30630-7)
- Barabási A-L. 2016 *Network science*. Cambridge, UK: Cambridge University Press.
- Bonell C, Michie S, Reicher S, West R, Bear L, Yardley L, Curtis V, Amlôt R, Rubin GJ. 2020 Harnessing behavioural science in public health campaigns to maintain 'social distancing' in response to the covid-19 pandemic: key principles. *J. Epidemiol. Community Health*. **74**, 617–619.
- Chater AM, Arden M, Armitage C. 2020 *Behavioural science and disease prevention: psychological guidance*. Leicester, UK: The British Psychological Society.
- West R, Michie S, Rubin GJ, Amlôt R. 2020 Applying principles of behaviour change to reduce SARS-CoV-2 transmission. *Nat. Human Behav.* **4**, 451–459. (doi:10.1038/s41562-020-0887-9)
- Fenichel EP. 2013 Economic considerations for social distancing and behavioral based policies during an

- epidemic. *J. Health Econ.* **32**, 440–451. (doi:10.1016/j.jhealeco.2013.01.002)
27. Vlaev I, King D, Darzi A, Dolan P. 2019 Changing health behaviors using financial incentives: a review from behavioral economics. *BMC Public Health* **19**, 1–9. (doi:10.1186/s12889-019-7407-8)
 28. Independent SAGE. Towards an integrated find, test, trace, isolate, support (FTTIS) response to the pandemic, 2020. See www.independentsage.org/wp-content/uploads/2020/06/IndependentSAGE-report-4.pdf.
 29. Fehr E, Falk A. 2002 Psychological foundations of incentives. *Eur. Econ. Rev.* **46**, 687–724. (doi:10.1016/S0014-2921(01)00208-2)
 30. Steptoe A, Willemsen G. 2004 The influence of low job control on ambulatory blood pressure and perceived stress over the working day in men and women from the Whitehall II cohort. *J. Hypertens.* **22**, 915–920. (doi:10.1097/00004872-200405000-00012)
 31. Webster RK, Brooks SK, Smith LE, Woodland L, Wessely S, Rubin GJ. 2020 How to improve adherence with quarantine: rapid review of the evidence. *Public Health* **182**, 163–169. (doi:10.1016/j.puhe.2020.03.007)