

ARTICLE

ENDOGENOUS SOCIAL DISTANCING AND CONTAINMENT POLICIES IN SOCIAL NETWORKS

Fabrizio Adriani* and Dan Ladley

University of Leicester School of Business, University of Leicester, Leicester, United Kingdom

*Corresponding author. Email: fa148@le.ac.uk

(Received 9 April 2021; revised 18 June 2021; accepted 30 June 2021)

Can smart containment policies crowd out private efforts at social distancing? We analyse this question from the perspective of network formation theory. We focus in particular on the role of externalities in social distancing choices. We also look at how these choices are affected by factors such as the agents' risk perception, the speed of the policy intervention, the structure of the underlying network and the presence of strategic complementarities. We argue that crowding out is a problem when the probability that an outbreak may spread undetected is relatively high (either because testing is too infrequent or because tests are highly inaccurate). This is also the case where the choice of relaxing social distancing generates the largest negative externalities. Simulations on a real-world network suggest that crowding out is more likely to occur when, in the absence of interventions, face-to-face contacts are perceived to carry relatively high risk.

Keywords: social distance; networks; test & trace; contagion; crowding out.

JEL codes: D85; I12; D62.

1. Introduction

In the absence of pharmaceutical interventions, containment is often the only tool available to policy-makers seeking to mitigate the consequences of an epidemic. Nonetheless, indiscriminate lockdowns and quarantines carry huge human and economic costs. In the immediate aftermath of the initial COVID-19 outbreak, these considerations have led the search for alternatives to total lockdowns. Across the world, nearly all advanced economies have adopted, in different measures and with varying results, a combination of testing and contact tracing to detect and isolate outbreaks. These measures are usually combined with localised interventions (school closures, event cancellations and quarantines) following the detection of an outbreak. In this context, a concern shared by many behavioural scientists is the extent to which these smart approaches to containment might displace personal efforts at social distancing, a phenomenon reminiscent of 'crowding out' results familiar from macroeconomics. The idea behind crowding out is that, if people perceive that containment is doing its job, they could relax social distancing and enjoy more face-to-face contacts. However, as everybody does this, the risk of infection may increase again. These behavioural responses are relatively difficult to measure empirically because of the lack of a clear counterfactual. In cases like this, theory can help to understand some qualitative features of the problem and narrow down the set of plausible scenarios. In this paper, we use strategic network formation theory to analyse the effect of plausible responses to containment. In order to better illustrate the mechanisms at work, we sketch a simple static model where agents optimally choose the degree of social distancing they wish to keep. Consistent with a long-established tradition in economics, we thus analyse how agents change their behaviour in response to policy interventions aimed at reducing risk. More precisely, we assume that agents are part of an exogenously given social structure,

which may reflect economic, social, geographical or regulatory constraints. Within this structure, agents choose which of their links they wish to keep active. Active links provide monetary or nonmonetary benefits, but increase the probability of contagion. Hence, agents may want to render some of their links inactive to reduce risk (social distancing). For instance, people may want to refrain from meeting some of their friends, colleagues or extended family in person and interact with them online. Adriani and Ladley (2021) formally analyse the model, and provide several analytical results and a number of insights from simulations. Here, we informally review some of the results, analyse further implications and explore some plausible counterfactuals by running simulations on a real-world network. In particular, we use data from the BBC Four Pandemic project to construct simulated networks and analyse how the spread of the disease is affected by policy interventions.

Specifically, the theory generates a number of insights that can help to answer the following questions:

1. When is crowding out an issue?
2. Can the intervention crowd out social distancing to the extent that the infection rate is actually higher than under no intervention?
3. How fast should the intervention be?
4. Does the shape of the social network matter?

Previewing some of the results, the answer to the first question is rather intuitive: crowding out only matters when chances that an outbreak will pass undetected are relatively high. This may, for instance, be the case when testing is imperfect, either because agents are tested infrequently or because the test itself is inaccurate. For example, evidence suggests that lateral flow tests (LFTs), which provide quick results and are usually deployed for mass asymptomatic testing, tend to be far less accurate than polymerase chain reaction tests.¹ Accordingly, practitioners worry that LFTs may provide false reassurances to tested individuals, which could potentially lead to an increased in overall risk. Our analysis suggests that a sufficiently imperfect test, if not coupled with strict mandatory social distancing rules, may indeed end up increasing the infection rate and reduce welfare.

This leads us straight into the answer to the second question. At first glance, the idea that the behavioural responses to an intervention could actually lead to a higher infection rate might appear in conflict with standard rational expectations arguments. After all, one is only going to relax social distancing if the intervention reduces risk. If the intervention were to increase the infection probability, people would be reluctant to resume face-to-face contacts. But then, according to this argument, crowding out cannot possibly occur in the first place. The problem with this intuition is that it does not take into account the role of externalities. When two agents agree to meet in person, they may not fully internalise the negative externality they exert on their other contacts. As a result, the intervention may trigger behavioural changes that, while individually optimal, may induce an increase in aggregate risk. As we argue below, this problem is more severe when strategic complementarities are present.

Simulation results suggest that the agents' risk perceptions play a crucial role in determining whether the infection rate may increase as a consequence of the intervention. There is more scope for crowding out when, in the absence of intervention, social contacts are perceived to carry high risk. The intuition is as follows. As active links are perceived to be riskier, the marginal increase in risk from an additional active link becomes sharply decreasing in the number of links one already has. Intuitively, if the chances to catch the disease from one's existing connections are already large, an additional connection does not add much risk. In these circumstances, it is usually optimal to have either very few or a large number of active links, but not a moderate number. This makes the optimal number of active links relatively elastic with respect to policy interventions. Hence, even a very inaccurate intervention may induce agents to significantly reduce social distancing, which in turn increases the infection rate.

Question 3 is the main focus of Adriani (2020) and Adriani and Ladley (2021). Keeping agents' behaviour fixed, it is obviously optimal to stop the spread of the disease as soon as a case is detected.

¹Source: *The British Medical Journal* news, doi: <https://doi.org/10.1136/bmj.m4848>.

However, things change when social distancing is endogenous. Ex-ante, ‘fast’ interventions—which aim at preventing any transmission to neighbours as soon as an outbreak is detected—are more effective when accuracy is high. In contrast, when accuracy is low, ‘slow’ interventions—that allow a limited initial spread of the outbreak—may sometimes perform better. Intuitively, a slow intervention does not protect much those who choose high exposure, but may protect some of their contacts, since there is a chance that the outbreak might be contained before it reaches them. In the context of the model, a slow intervention does not provide much incentive to relax social distancing and, at the same time, it is effective in limiting the extent of the negative externality provided by those who relax it. The drawback of course is that an intervention that is too slow may fail to meaningfully contain the spread of the disease, especially when the social network is very dense.

This in turn leads us to the last question. The worst outbreaks of COVID-19 have taken place in densely populated urban areas like New York and London, and the large majority of cases can be traced to relatively few ‘superspreaders’. This suggests the shape of the underlying network may be crucial for shaping optimal containment policies. Simulated results show that indeed higher network density facilitates the spread of the virus. However, the increase in infection rate appears to be uniform with respect to the accuracy of the intervention and higher network density does not qualitatively affect the extent of crowding out.

The paper is organised as follows. After a brief review of the literature, we sketch the model and provide a number of theoretical insights in Section 2. Section 3 focuses on the simulation results. Concluding remarks are in Section 4.

1.1. Related literature

At least since Peltzman (1975), economists have been interested in whether behavioural responses to risk reducing interventions may offset the intervention’s direct effect.² Following Kremer’s (1996) seminal work, some of the recent works specifically focus on infection prevention and mitigation. These include Geoffard and Philipson (1996), Gersovitz and Hammer (2004), Goyal and Vigier (2015), Greenwood *et al.* (2019), Reluga (2010), Rowthorn and Toxvaerd (2012), and Toxvaerd (2020). These models, however, tend to abstract from social structure.³

The epidemiological literature on contagion in social networks is too vast to be reviewed here (see Keeling and Eames, 2005, for a review). We thus mainly restrict attention to the Economics literature. Examples include Acemoglu *et al.* (2016a), Cabrales *et al.* (2012), Erol and Vohra (2018), Goyal *et al.* (2016), and Goyal and Vigier (2014). These works tend to mostly focus on network design and/or resilience and mostly abstract from specific policy interventions (notable exceptions are Demange, 2018, and Galeotti and Rogers, 2013). A similar point broadly applies to the literature on financial contagion (see the surveys in Acemoglu *et al.*, 2016b, and Glasserman and Young, 2016).

More similar to us, Blume *et al.* (2013) analyse welfare in a static, reduced-form model of contagion in networks (see also Bougheas, 2018). Decentralised equilibria are suboptimal because of the externality outlined in the introduction: when forming links, agents do not internalise the increase in the probability of infection of their neighbours. Using a variant of Blume *et al.* (2013), Talamàs and Vohra (2020) study the welfare consequences of the introduction of an imperfect vaccine. They point out how a sufficiently imperfect vaccine may increase the risk of contagion and reduce welfare. In the context of testing, Acemoglu *et al.* (2020a) and Adriani (2020) both study crowding out under an imperfect testing technology. The settings and the mechanisms at work are, however, quite different. Acemoglu *et al.* (2020b) consider a setting with strategic substitutabilities and benefits from social interaction linear in the level of social activity. This implies that, in

²See also Hoy and Polborn (2015).

³The current pandemic has prompted a number of works on social distancing and interventions (e.g., Acemoglu *et al.*, 2020b; Alvarez *et al.*, 2020; Galanis, 2021; Kruse and Strack, 2020; Makris, 2021; Muscillo *et al.*, 2020), which, however, do not specifically focus on offsetting behaviour.

equilibrium, agents randomise between perfect isolation and full exposure. In contrast, Adriani (2020) considers a network formation game which displays strategic complementarities and assumes decreasing returns from active links, so that agents tend to choose intermediate levels of social activity.

2. Theoretical framework

2.1. A static model of contagion with endogenous network formation

In this section, we sketch the baseline model—full details are available in Adriani and Ladley (2021). There are $N \geq 2$ nodes organised on a graph (the *underlying network*). This network summarises the links available to each node and reflects the contacts that people would normally have. During an epidemic, each node can choose to make any of these links active or to keep them inactive. Only links featured in the underlying network can be activated. This is meant to capture the presence of an exogenous social structure generated by social or geographical constraints. On top of this structure, agents choose their degree of social distancing (how many links they wish to keep active). If node i already has $n \geq 0$ active links, an additional active link generates a benefit $b\delta^{n+1}$, with $b > 0$ and $\delta \in (0, 1)$. Since δ lies between 0 and 1, there are decreasing returns from active links. These may reflect time constraints, the fact that some connections are more important/valuable than others, or other economic considerations. After nodes have decided which links they want to keep active, one of the N nodes (the *initial carrier*) becomes infected with exogenous probability θN , $\theta \in (0, 1/N]$. If a node is infected, the disease can spread through the active links. More precisely, we assume that each node is immune with some probability ϕ . An immune node is protected from infection and does not infect its neighbours. A node is therefore infected if and only if it is either the initial carrier or it is not immune and there is a path of nonimmune nodes connecting it to the initial carrier. We assume nodes do not know whether they are immune or not.⁴ Finally, nodes that become infected suffer a loss $L > 0$. The model's parameters are summarised in table 1 (see Section 3).

Table 1. Model's parameters

Parameters	Values	Description
N	411	Number of nodes
θ	0.0124	Exogenous probability of infection
ϕ	0, 0.1, 0.2, ..., 0.5	Share of immune nodes
L	2045, 1700, 1400	Expected utility loss from infection
b	314	Utility from active links
δ	0.765	Rate of decay of utility from add. link
λ	0, 0.1, 0.2, ..., 1	Interv. accuracy (Pr outbreak detected)
$\bar{\pi}$	0.0124, 0.0206, 0.0412, ...	Perceived prevalence
	..., 0.0824, 0.1648	(Pr neighbour is infectious)

⁴This would be the case if immunity were determined by genetic traits that tend to vary in the population. Alternatively, the same would apply if, during previous outbreaks, a large number of infections were asymptomatic and/or tests were initially restricted to clinically severe cases, as in Spring 2020. For instance, Verity *et al.* (2020) estimate that a large proportion of COVID-19 cases were not identified.

2.1.1. Intervention

Under an intervention, nodes are tested. If there is an outbreak, the initial carrier is detected and isolated with probability λ . If the initial carrier is isolated, the infection does not spread to other nodes. We refer to λ as the accuracy of the intervention. This reflects how systematically nodes are tested or how accurate the test is. We will discuss, in some cases, the effect of ‘slow’ interventions, where the initial carrier has time to infect her first-degree neighbours (provided they are not immune) before the outbreak is contained.

2.1.2. Information

Adriani and Ladley (2021) consider two versions of this model, which differ in the degree of agents’ sophistication. The first version analyses the case where agents are sophisticated and have full knowledge of the underlying network. At the other extreme, we consider the case where agents only have local knowledge of the underlying network, in the sense that they know what links are available to them but ignore what links are available to other agents.

2.1.3. Equilibrium

The equilibrium concept we use for the full information version of the model is pairwise stability (Jackson and Wolinsky, 1996). In a pairwise stable network, no two agents could benefit from activating a link and no agent could profit from unilaterally cutting one of their link.⁵ When agents only have local knowledge of the network, it is natural to assume that each agent chooses a target number of links that they want to keep active. An equilibrium is a profile of target numbers such that each agent’s target is optimal given the strategy of other agents. Available links are then activated in random order until each agent has either reached his/her target or has no available link to activate. This version of the model will be used in the simulations.

2.2. Externalities, complementarity and substitutability

We illustrate the key mechanisms at work by considering a simple example with $N = 3$. In the figure below, suppose that Alice and Bob meet frequently (the solid line). For instance, they may be co-workers. Bob and Charles are close friends, but they are currently only talking via Skype because of social distancing (the dashed line).

Consider Bob’s incentive to activate his link to Charles. In the absence of any intervention, Bob’s chances to be infected by Charles are $(1 - \phi)\theta$, that is, the probability that Bob is not immune times the probability that Charles is a carrier. Since Bob already has an active link, his benefit from activating the additional link to Charles is $b\delta^2$. If $b\delta^2 < (1 - \phi)\theta L$, then there is no incentive for Bob to activate the link. Suppose now that Charles is tested but the test generates a false negative with probability $1 - \lambda$. In this case, Bob’s probability to be infected by Charles drops to $(1 - \lambda)(1 - \phi)\theta$. If

$$(1 - \lambda)(1 - \phi)\theta L < b\delta^2 < (1 - \phi)\theta L, \quad (1)$$

then the intervention will induce Bob to activate the link. This is, however, not the only effect. Bob’s decision provides a negative externality to Alice. In the absence of intervention, Alice’s probability to be infected is $\theta + (1 - \phi)\theta$, namely the probability that she is the carrier plus the probability to catch the disease from Bob. After the intervention, her probability of infection increases by $(1 - \lambda)(1 - \phi)^2\theta$, namely the probability to be infected by Charles via Bob.⁶ As a result of the externality, it may thus be

⁵This version of the model is fully analysed in Adriani and Ladley (2021).

⁶This only happens when both Alice and Bob are not immune, which occurs with probability $(1 - \phi)^2$.

that, even if the meetings between Bob and Charles have indeed become less risky if taken in isolation, the overall risk of infection has increased.⁷

The story does not end here, though. Suppose that Alice is also very fond of Charles. Upon hearing that Bob and Charles regularly go out for a drink, she will probably decide that it does not make sense for her to avoid meeting Charles while still bearing the risk of being infected by Charles via Bob. Hence, she may decide to start meeting Charles as well. Indeed, if the link between Bob and Charles is already active, activating a direct link to Charles increases Alice's risk only in the event that Bob is immune (so that she cannot catch the infection via Bob). Formally, when the link between Bob and Charles is active, her probability of infection increases only by $(1 - \lambda)(1 - \phi)\phi\theta$ when activating her link to Charles, compared to $(1 - \lambda)(1 - \phi)\theta$ when the link between Bob and Charles is inactive. Since Alice has stronger incentive to activate her link to Charles when the Bob–Charles link is active, Bob's and Alice's propensities to meet Charles are *strategic complements*. In a nutshell, if everyone you are in contact with has high levels of social activity, only two options make sense. Either you cut all links (if you can afford it), or you start partying as well. In the presence of externalities and complementarities, it is possible that the intervention may actually generate a higher infection rate and lower welfare.

This result crucially rests on the testing technology being inaccurate. Indeed, it is always possible to find (sufficiently high) values of λ such that the intervention reduces the overall probability of infection.⁸ Intuitively, even under maximum crowding out (e.g., agents choose to activate all their available links), if we could immediately spot anyone who carries the virus, contagion would be negligible. This suggests that crowding out is mostly a problem for low to intermediate levels of accuracy. Outbreaks must be spotted frequently enough to induce agents to reduce social distancing but not so frequently that the direct effect of the intervention in reducing risk outweighs the indirect crowding out effect.

This is not meant to imply that mass testing should be avoided if the available test is inaccurate. However, it casts doubts on the extent to which mass testing can be used to ease mandatory social distancing restrictions. Indeed, mass testing may increase the need for stricter restrictions.

2.3. Speed of intervention

Experts and policymakers worry that slow interventions may fall 'one step behind' the outbreak. If a carrier is not identified and notified quickly, they may go on to infect others who then may infect others and so on. In the absence of behavioural responses, it is indeed always better to have interventions that are as fast as possible. This is witnessed by the scramble to find technological solutions that speed up contact tracing. However, when we allow for behavioural responses, the picture becomes somehow murkier. In particular, there is a complementarity between the accuracy and the speed of intervention. Inaccurate interventions that are too fast may be counterproductive. This phenomenon is explored in detail in Adriani and Ladley (2021). The logic of the result is that a slightly slower intervention may help to weaken both the effect of the externality and strategic complementarities. To see this, consider again [figure 1](#). Suppose that contact tracing is not fast enough to prevent Bob from being infected if Charles is infected but sufficiently fast to spare Alice so long as she does not enter in direct contact with Charles. This has two effects. First, it reduces Bob's incentive to resume his meetings with Charles in the first place,

⁷Formally, under no intervention, Charles has probability θ to catch the disease. Alice and Bob have both probability $\theta + (1 - \phi)\theta$. Hence, the average infection probability is

$$\theta + \frac{2}{3}(1 - \phi)\theta. \quad (2)$$

If an intervention with accuracy λ induces Bob and Charles to activate their link, the average infection probability becomes

$$\theta + \frac{4}{3}(1 - \lambda)(1 - \phi)\theta + \frac{2}{3}(1 - \lambda)(1 - \phi)^2\theta. \quad (3)$$

Hence, if $\lambda < (2 - \phi)/(3 - \phi)$, the average probability increases.

⁸Formally, there always exists $\lambda^* < 1$ such that the intervention reduces the overall infection probability for all $\lambda < \lambda^*$.





Figure 1. A network with three nodes. Solid lines indicate active links. Dashed lines inactive links

since he knows that he has a fair chance to catch the virus if he does. Second, if Bob and Charles decide to go on and meet anyway, it reduces Alice's incentive to join them, since there is a good chance that a possible outbreak may be contained before it reaches her—so long as she does not have direct contacts with Charles. Another way to interpret this result is that it is unlikely that one needs to worry both about containment falling one step behind the spread of the virus and about crowding out. One tends to exclude the other.

2.4. Shape of the underlying network

We now consider how the results depend on the shape of the network. In dense networks, where everyone constantly meets everyone else, the risk of falling one step behind the outbreak is a real one. Moreover, if there are individuals who are central in the network or who tend to meet a lot of other people, they should be tested with high frequency. For instance, consider the following three networks:



The network on panel (a) is dense, and any outbreak is likely to spread very quickly to all nodes that are not immune. The network in the middle is relatively sparse. It is sufficient that Bob or Charles is immune (or that the outbreak is contained by the time it reaches them) to stop contagion from spreading further. In panel (c), the network is again sparse, but contagion may hit most nodes if Alice is infected. In panel (a), it seems natural to test everyone with the same frequency. In panel (c), a more sensible approach would be to concentrate testing on Alice, since any infection spreading to the others must pass through her. For instance, so long as visits are not allowed, staff are the main gateway for outbreaks hitting care homes. It thus makes sense to test them frequently. Remember, however, that the shape of the network is not independent of the approach to containment. For instance, while in panel (a), an outbreak is likely to spread very quickly, the likelihood that networks of such density may form depends on the type of intervention in place. In what follows, it will be important to distinguish between the network's *underlying density*—that is, the density of the network formed by the available links—and its *actual density*, which refers to the portion of the network that is active during an epidemic.

3. Simulations

The version of the model with fully rational agents who perfectly know the underlying network is computationally demanding and unrealistic for large networks. We thus use the version with limited knowledge of the network to simulate behavioural responses to an intervention. Moreover, we do not constrain agents to have rational expectations, but instead introduce an additional expectation parameter, $\bar{\pi}$, capturing the perceived prevalence of the virus among the population. More precisely, $\bar{\pi}$ is the probability that each node assigns to any of their neighbours carrying the virus in the absence of any intervention. The agent's perceived risk is thus determined by two parameters: $\bar{\pi}$ and λ (the intervention's accuracy). We allow for multiple values of both parameters to see how behaviour depends on agents' beliefs. Clearly enough, anchoring agents' expectations on a fixed parameter instead of equilibrium behaviour comes at the cost of ignoring potential strategic interactions. However, deriving probability distributions over the set of all possible active networks resulting from the strategies of a large number of agents is computationally daunting. The approach we follow places far less demands on the mental capacity of agents.⁹ The model used for the simulations is explained in detail in the Appendix.

The underlying network is built using data from the BBC Four Pandemic project collected in 2017 in the Surrey town of Haslemere. The data track the contacts (physical distance) of 469 participants in the project over 3 days. As the data were collected before the Covid pandemic, we interpret the observed contacts as the links that would be available to agents in normal circumstances (i.e., with a zero infection rate). We restrict attention to contacts within 2- or 4-m distance. This leaves us with 392 connected nodes with degrees ranging from 1 to 37 links (average 2.77). Given this underlying network, we use the model to determine how many of these links agents may want to keep active under various configurations of the model's parameters. The values chosen for the parameters are displayed in table 1. Details on how the parameters' values were chosen and the methodology used in the simulations are available in the Appendix. Here, it is worth noting that, inevitably, there is high uncertainty about the value of some of the parameters.¹⁰

Although the BBC Four Pandemic dataset has been previously used by epidemiologists to study disease transmission in networks (see Firth *et al.*, 2020a, 2020b; Kissler *et al.*, 2020; Klepac *et al.*, 2018), the data have several limitations. This calls for some caution in interpreting our results. We see our simulations more like accounts of plausible counterfactuals rather than accurate measurements. A potential problem is that we observe only a portion of the whole network available to participants (i.e., only those who volunteered to take part in the project), which covers only a fraction of the town's population. This results in a relatively sparse underlying network if we restrict attention to contacts that are within 2 m of distance. To partially compensate for this, we look at what happens when we vary the underlying density by allowing for more or less physical distance. Another question is, of course, the extent to which volunteers who participated in the project are representative of the population at large.¹¹

Figure 2 shows the share of infected agents conditional on an outbreak occurring (solid line) for multiple values of the intervention's accuracy λ . The dashed line depicts the optimal number of active links for each agent. We start with the case where agents perceive risk to be relatively low and then gradually increase the value of $\bar{\pi}$. The $\lambda = 0$ case returns the conditional infection rate under no intervention. For low values of $\bar{\pi}$ (the top panels), the infection rate monotonically decreases with the accuracy of the intervention. This implies that any intervention outperforms no intervention independently of its accuracy, so that there is no crowding out. Things, however, change when agents assess risk to be relatively high (the two bottom panels). In these cases, at low levels of accuracy, the intervention increases the infection rate. The rationale for the relationship between risk perceptions and crowding out

⁹Adriani and Ladley (2021) characterise the equilibrium and provide comparative statics on the intervention accuracy for a general version of the model which nests both this case and the rational expectations case.

¹⁰Moreover, while we take a representative agent approach for simplicity, some of the parameters' values are likely to differ across individuals.

¹¹For instance, children under the age of 13 were not included in the project, and ownership of a smartphone—which was necessary to participate—tends to be less frequent among over 65.

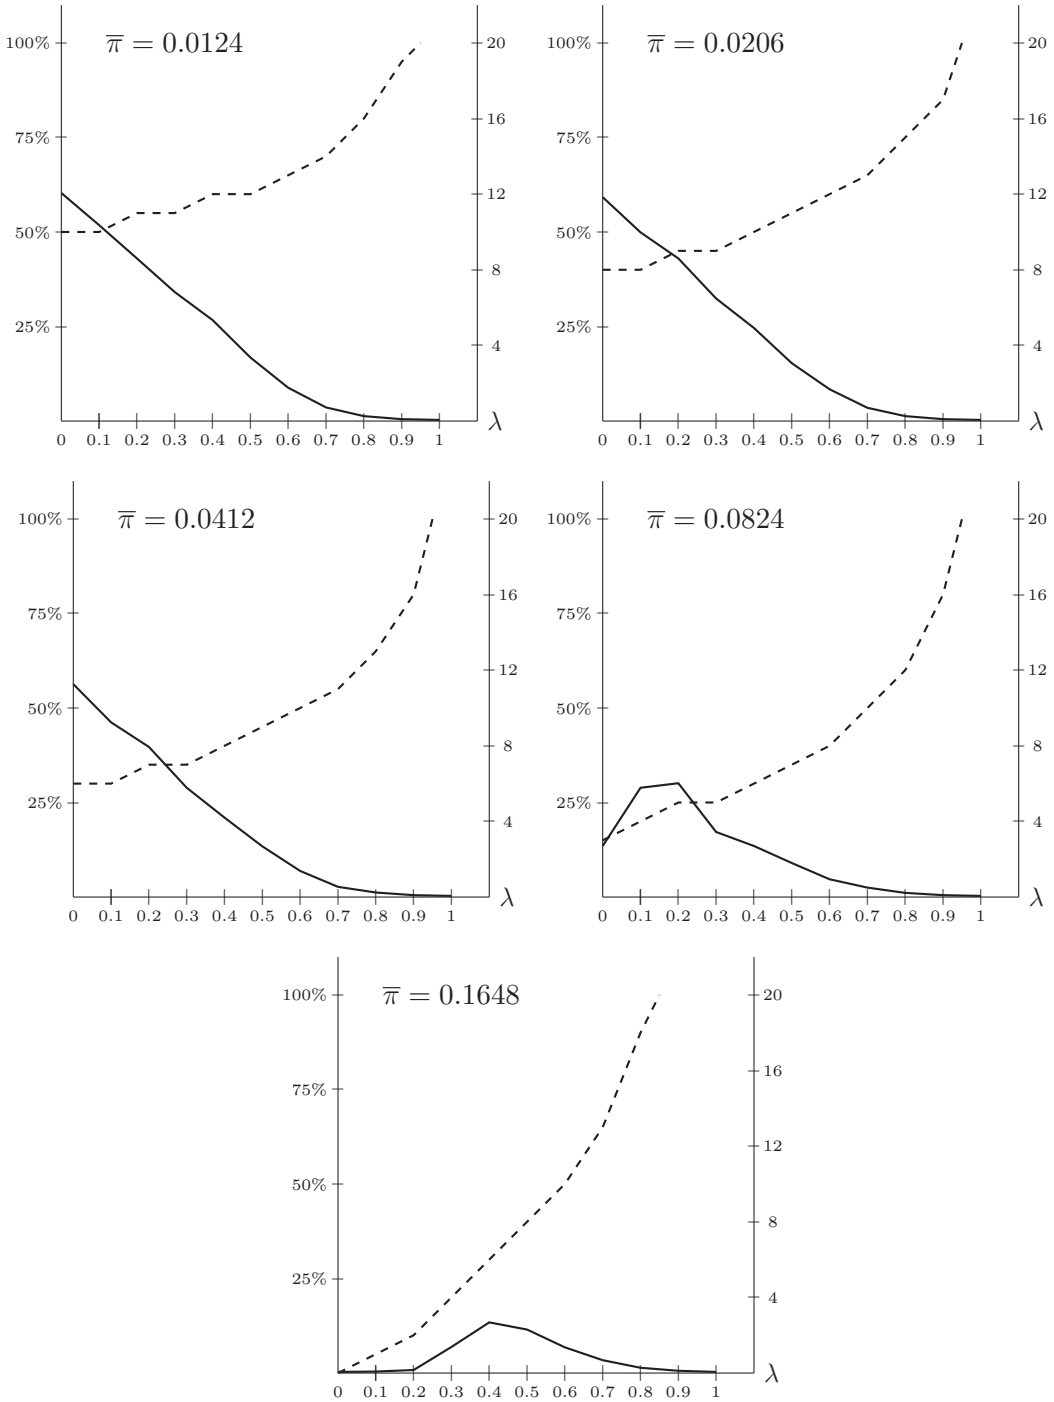


Figure 2. Relation between share of infected agents (left axis, solid line) and intervention accuracy, λ , for different levels of perceived risk $\bar{\pi}$. The dashed line measures the target number of active links (right axis). Parameters: $\phi = 0.3$, $\bar{\pi} = 0.0824$, $L = 2045$, $\delta = 0.765$, $b = 314$. Links are contacts within a physical distance of 2 m. The share of infected agents is conditional on an outbreak hitting the network

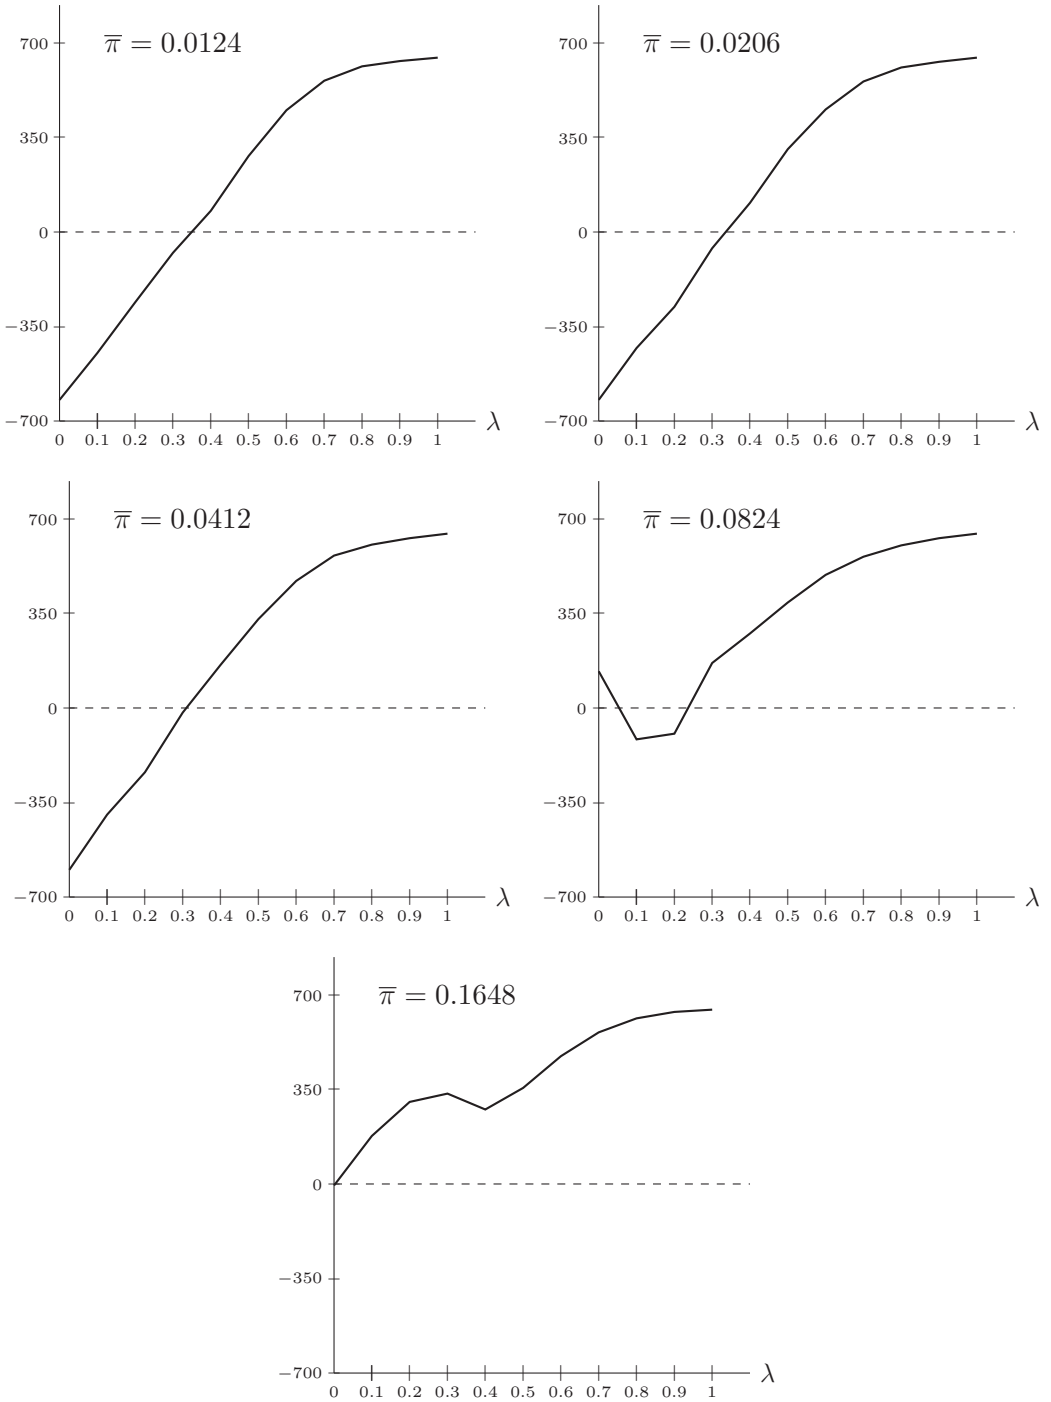


Figure 3. Relationship between welfare (conditional on an outbreak hitting the network) and intervention accuracy, λ , for different levels of perceived risk $\bar{\pi}$. Parameters: $\phi = 0.3$, $\bar{\pi} = 0.0824$, $L = 2045$, $\delta = 0.765$, $b = 314$. Links are contacts within a physical distance of 2 m

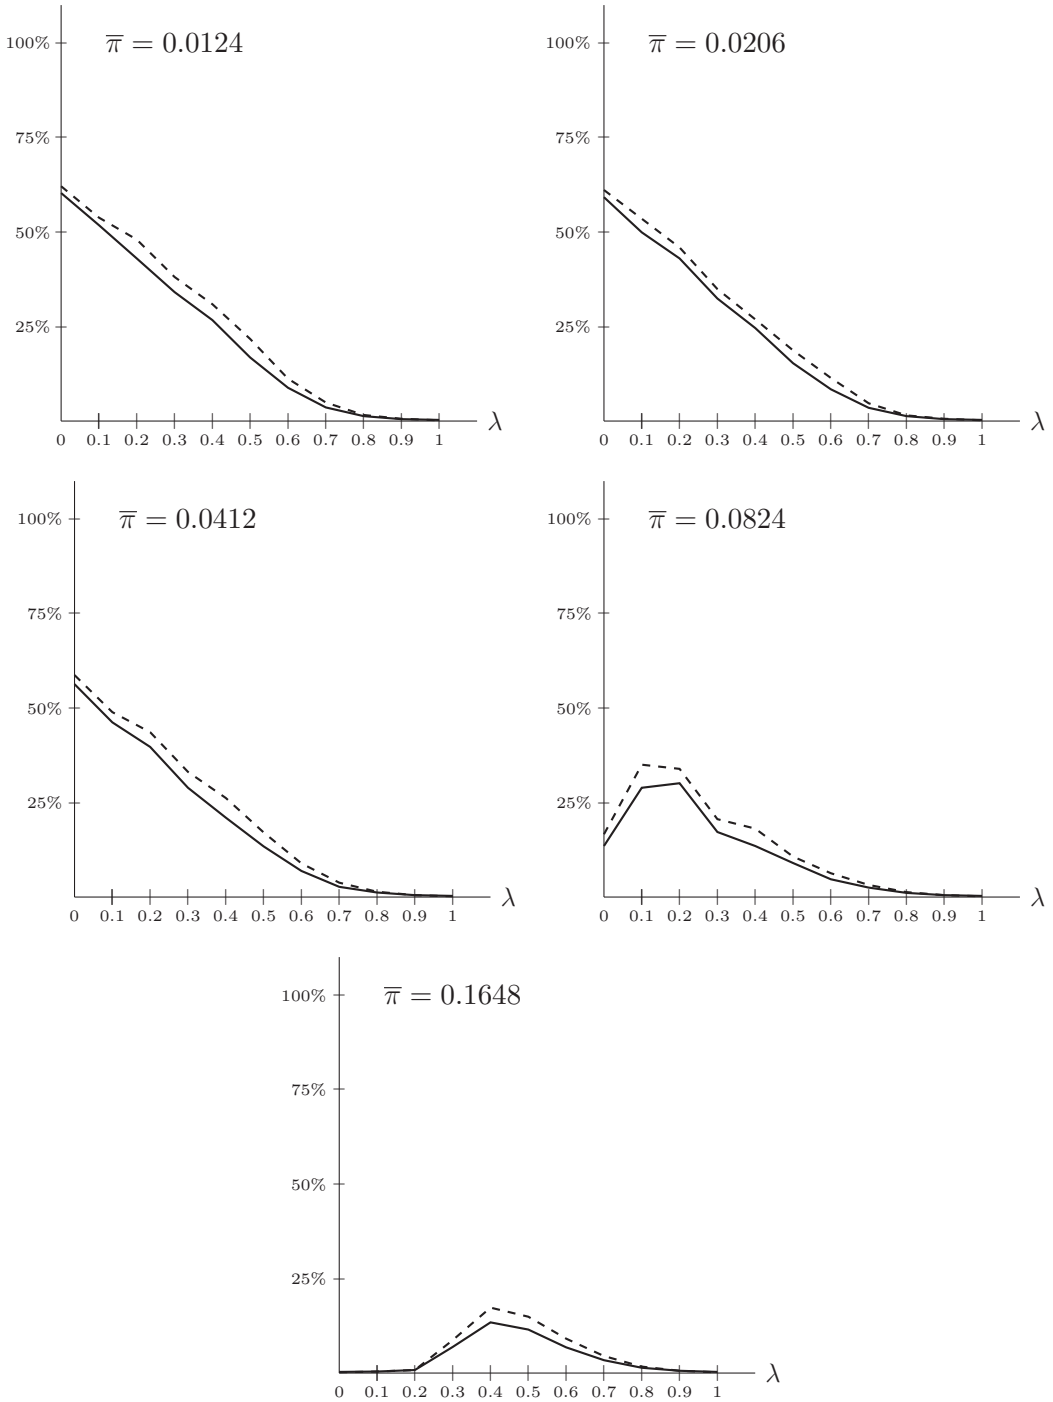


Figure 4. Relation between share of infected agents and intervention accuracy, λ , for different levels of perceived risk $\bar{\pi}$ and different densities of the underlying network. The solid line refers to the network obtained by restricting attention to physical distances within 2 m. The dashed line refers to the denser network obtained with distances within 4 m. Parameters: $\phi = 0.3$, $\bar{\pi} = 0.0824$, $L = 2045$, $\delta = 0.765$. The share of infected agents is conditional on an outbreak hitting the network

is the following. When face-to-face meetings are perceived to carry high risk, the marginal increase in risk from a second active link is quite smaller than the marginal increase from the first, since one can only be infected once. In turn, the marginal increase from the third is much smaller than the second and so on. In this situation, it is usually optimal to either have very few active links or to have many, but not an intermediate number. This makes the desired number of active links very elastic with respect to the policy intervention. Intuitively, even a very inaccurate intervention may be enough to induce agents to switch from meeting almost no one to meeting a lot of people. Clearly enough, such an environment is highly susceptible to crowding out. It is also worth noting that, while the results provided in figure 2 are for a specific configuration of parameters, the relationship between $\bar{\pi}$ and crowding out is robust to using different parameter values.

Looking at the infection rate provides an incomplete picture, since it takes into account the costs but not the benefits of social activity. In order to also account for the benefits of active links, we consider agents' welfare in figure 3. To measure welfare, we take the average ex-post utility across agents and repetitions conditional on an outbreak hitting the network. The pattern observed is consistent with the results on the infection rate. For low perceived risk, welfare monotonically increases in the accuracy of the intervention. In contrast, when risk is perceived to be high, we observe crowding out and, more generally, nonmonotonicity.

Figure 4 considers what happens when we change the network's underlying density by allowing contacts within a 4-m physical distance.¹² This increases the average number of available links from 2.77 to 3.06. It is, however, worth emphasising again that an increase in the average degree does not mechanically translate into an increase in risk as agents can choose whether to activate or not these links. In other words, higher underlying density does not necessarily imply higher density in the active portion of the network. The solid line shows the effect of the intervention for the 2-m case, whereas the dashed line depicts the effect for the denser network obtained in the 4-m case. As expected, higher density increases the conditional infection rate. The effect of density is, however, pretty much uniform and does not seem to interact with crowding out. The reason is that, as already explained, crowding out tends to occur when agents perceive risk to be high, so that they tend to keep active only a subset of the available links. In this situation, making more links available to them only has limited effects.

4. Concluding remarks

Network formation theory suggests that, in the presence of social distancing externalities, interventions aimed at detecting and containing outbreaks may crowd out social distancing and end up increasing the infection rate. Our results indicate that this is likely to happen only when interventions tend to miss most outbreaks, either because the testing technology is inaccurate or because of limited coverage. We also found that crowding out tends to occur when, under no intervention, agents are reluctant to engage with others because of high perceived risk. In this case, the marginal increase in the probability of infection from an additional contact decreases sharply in the number of contacts one already has, so that agents' optimal number of contacts become very elastic to policy changes. Further research is required to determine whether these hypotheses have empirical support.

On the theoretical side, the analysis can be extended in several directions. First, we did not take into account other regarding preferences. While these may appear at first sight as a way to redress/mitigate the externalities, recent theoretical results suggest that their effect is more complex (see Toxvaerd, 2021). Second, we essentially focused on a representative agent. However, agents are likely to differ in several of their characteristics. For example, Acemoglu *et al.* (2020b) consider the interaction between two groups (high and low risk) differing in their preferences. Empirical evidence suggests that differences in beliefs may also play an important role (see Bordalo *et al.*, 2020). Introducing heterogeneity of beliefs about the prevalence of the disease or its potential harm may thus provide another interesting avenue for research.

¹²Firth *et al.* (2020b) also use a 4-m distance.

References

- Acemoglu, D., Chernozhukov, V., Werning, I. and Whinston, M.D.** (2020a), 'A multi-risk SIR model with optimally targeted lockdown', Working paper 27102, National Bureau of Economic Research, Cambridge, Massachusetts.
- Acemoglu, D., Makhdoumi, A., Malekian, A. and Ozdaglar, A.E.** (2020b), 'Testing, voluntary social distancing and the spread of an infection', Working paper 27483, National Bureau of Economic Research, Cambridge, Massachusetts.
- Acemoglu, D., Malekian, A. and Ozdaglar, A.** (2016a), 'Network security and contagion', *Journal of Economic Theory*, **166**, pp. 536–85.
- Acemoglu, D., Ozdaglar, A. and Tahbaz-Salehi, A.** (2016b), 'Networks, shocks, and systemic risk', in Bramoullé, Y., Galeotti, A. and Rogers, B. (eds), *The Oxford Handbook of the Economics of Networks*, Oxford: Oxford University Press, pp. 569–610.
- Adriani, F.** (2020), 'Social distance, speed of containment and crowding in/out in a network model of contagion', CeDeX Discussion Paper 2020-10.
- Adriani, F. and Ladley, D.** (2021), 'Social distance, speed of containment and crowding in/out in a network model of contagion', *Journal of Economic Behavior and Organization*, forthcoming.
- Alvarez, F.E., Argente, D. and Lippi, F.** (2020), 'A simple planning problem for Covid-19 lockdown', Working paper 26981, National Bureau of Economic Research, Cambridge, Massachusetts.
- Blume, L., Easley, D., Kleinberg, J., Kleinberg, R. and Tardos, É.** (2013), 'Network formation in the presence of contagious risk', *ACM Transactions on Economics and Computation (TEAC)*, **1**, 2, pp. 1–20.
- Bordalo, P., Coffman, K.B., Gennaioli, N. and Shleifer, A.** (2020), 'Older people are less pessimistic about the health risks of COVID-19', Working paper 27494, National Bureau of Economic Research, Cambridge, Massachusetts.
- Bougheas, S.** (2018), 'Contagion in networks: Stability and efficiency', Fondazione ENI Enrico Mattei working paper 050.2017.
- Cabrales, G., Gottardi, P. and Vega-Redondo, F.** (2012), 'Risk-sharing and contagion in networks', Working paper 2013/01. European University Institute, Florence.
- Demange, G.** (2018), 'Contagion in financial networks: A threat index', *Management Science*, **64**, 2, pp. 955–70.
- Emery, J.C., Russell, T.W., Liu, Y., et al.** (2020), 'The contribution of asymptomatic SARS-CoV-2 infections to transmission on the Diamond Princess cruise ship', *Elife*, **9**, p. e58699.
- Erol, S. and Vohra, R.** (2018), 'Network formation and systemic risk', SSRN 2546310.
- Firth, J.A., Hellewell, J., Klepac, P., Kissler, S., Kucharski, A.J. and Spurgin, L.G.** (2020b), 'Using a real-world network to model localized COVID-19 control strategies', *Nature Medicine*, **26**, 10, pp. 1616–22.
- Firth, J.A., Hellewell, J., Klepac, P., et al.** (2020a), 'Combining fine-scale social contact data with epidemic modelling reveals interactions between contact tracing, quarantine, testing and physical distancing for controlling COVID-19', medRxiv 2020.05.26.20113720.
- Galanis, S.** (2021), 'Social proximity as a tool to fight pandemics', *National Institute Economic Review*, **257**, pp. xx–xx.
- Galeotti, A. and Rogers, B.W.** (2013), 'Strategic immunization and group structure', *American Economic Journal: Microeconomics*, **5**, 2, pp. 1–32.
- Geoffard, P.Y. and Philipson, T.** (1996), 'Rational epidemics and their public control', *International Economic Review*, **37**, 3, pp. 603–24.
- Gersovitz, M. and Hammer, J.S.** (2004), 'The economical control of infectious diseases', *Economic Journal*, **114**, 492, pp. 1–27.
- Glasserman, P. and Young, H.P.** (2016), 'Contagion in financial networks', *Journal of Economic Literature*, **54**, 3, pp. 779–831.
- Goyal, S., Jabbari, S., Kearns, M., Khanna, S. and Morgenstern, J.** (2016), 'Strategic network formation with attack and immunization', in Cai, Y. and Vetta, A. (eds), *International Conference on Web and Internet Economics*, Berlin and Heidelberg: Springer, pp. 429–43.
- Goyal, S. and Vigier, A.** (2014), 'Attack, defence, and contagion in networks', *Review of Economic Studies*, **81**, 4, pp. 1518–42.
- Goyal, S. and Vigier, A.** (2015), 'Interaction, protection and epidemics', *Journal of Public Economics*, **125**, pp. 64–9.
- Greenwood, J., Kircher, P., Santos, C., and Tertilt, M.** (2019), 'An equilibrium model of the African HIV/AIDS epidemic', *Econometrica*, **87**, 4, pp. 1081–1113.
- Hoy, M. and Polborn, M.K.** (2015), 'The value of technology improvements in games with externalities: A fresh look at offsetting behavior', *Journal of Public Economics*, **131**, 1, pp. 12–20.
- Jackson, M.O. and Wolinsky, A.** (1996), 'A strategic model of social and economic networks', *Journal of Economic Theory*, **71**, 1, pp. 44–74.
- Keeling, M.J. and Eames, K.T.** (2005), 'Networks and epidemic models', *Journal of the Royal Society Interface*, **2**, 4, pp. 295–307.
- Kissler, S.M., Klepac, P., Tang, M., Conlan, A.J. and Gog, J.R.** (2020), 'Sparkling "The BBC Four Pandemic": Leveraging citizen science and mobile phones to model the spread of disease', bioRxiv, p. 479154.
- Klepac, P., Kissler, S. and Gog, J.** (2018), 'Contagion! The BBC Four Pandemic—The model behind the documentary', *Epidemics*, **24**, pp. 49–59.
- Kremer, M.** (1996), 'Integrating behavioral choice into epidemiological models of AIDS', *Quarterly Journal of Economics*, **111**, 2, pp. 549–73.
- Kruse, T. and Strack, P.** (2020), 'Optimal control of an epidemic through social distancing', Cowles Foundation Discussion Paper No. 2229.

- Makris, M. (2021), 'Covid and social distancing', *Economic Theory*, <https://doi.org/10.1007/s00199-021-01377-2>.
- Muscillo, A., Pin, P. and Razzolini, T. (2020), 'Covid19: Unless one gets everyone to act, policies may be ineffective or even backfire', *PLoS one*, **15**, 9, p. e0237057.
- Peltzman, S. (1975), 'The effects of automobile safety regulation', *Journal of Political Economy*, **83**, 4, pp. 677–725.
- Reluga, T.C. (2010), 'Game theory of social distancing in response to an epidemic', *PLoS Computational Biology*, **6**, 5, p. e1000793.
- Rowthorn, B.R. and Toxvaerd, F. (2012), 'The optimal control of infectious diseases via prevention and treatment', CEPR Discussion Paper No. DP8925.
- Smith, L.E., Amlôt, R., Lambert, H., et al. (2020), 'Factors associated with adherence to self-isolation and lockdown measures in the UK: A cross-sectional survey', *Public Health*, **187**, pp. 41–52.
- Talamàs, E. and Vohra, R. (2018), 'Go big or go home: Partially-effective vaccines can make everyone worse off', Working paper 18-006, Penn Institute for Economic Research, Philadelphia, Pennsylvania.
- Talamàs, E. and Vohra, R. (2020), 'Free and perfectly safe but only partially effective vaccines can harm everyone', *Games and Economic Behavior*, **122**, pp. 277–289.
- Toxvaerd, F. (2020), 'Equilibrium social distancing', Cambridge-INET Working Paper Series No. 2020/08.
- Toxvaerd, F. (2021), 'Contacts, altruism and competing externalities', CEPR Discussion Paper No. DP15903.
- Verity, R., Okell, L.C., Dorigatti, I., et al. (2020), 'Estimates of the severity of COVID-19 disease', medRxiv.

Appendix. Simulation methodology

In this section, we provide details for the simulations. Since the model used is a special case of Adriani and Ladley (2021), we refer the reader to the companion paper for further details.

A.1. Simulated model

We consider a network with $N > 1$ nodes indexed by $i = 1, \dots, N$. In the first stage, each node chooses a target number of active connections $n = 0, \dots, N - 1$ to maximise its expected payoff. The active network is then formed as follows. So long as there are at least two nodes with an available link who have not reached their target, a link between any two nodes who are below their target is randomly selected to be activated. The process stops when all nodes have either reached their target (or all their neighbours have reached their target) or have no available link left to activate.

We use a reduced form to model the perceived probability to be infected when activating a link. In particular, we assume that agents are totally unsophisticated and assign a fixed probability to be infectious, $\bar{\pi} \in (0, 1)$, to all other nodes. Under no intervention, the perceived probability to be infected for a node with n active links is thus

$$\theta + (1 - \phi)[1 - (1 - \bar{\pi})^n], \quad (4)$$

where the first term is the exogenous probability of infection and the second term is the probability that the node is not immune and at least one of its neighbours is infectious. In this particular setting, (4) is independent of the strategy of other nodes. See Adriani and Ladley (2021) for a general setting nesting a number of potential behavioural models, ranging from the case of unsophisticated agents to fully strategic agents with rational expectations. Consider now what happens when node i increases its target by one, from n to $n + 1$. When no new link becomes active (either because there are no additional links available to i or because all agents connected to i have reached their target), this is inconsequential. Conditional on an additional link becoming active, i 's expected payoff changes by

$$b\delta^{n+1} - \Delta_n L, \quad (5)$$

where $b\delta^{n+1}$ is the marginal benefit of an additional active link and Δ_n is the marginal increase in the perceived probability to be infected due to the new link,

$$\Delta_n = (1 - \phi) \left([1 - (1 - \bar{\pi})^{n+1}] - [1 - (1 - \bar{\pi})^n] \right) =$$

$$(1 - \phi)\bar{\pi}(1 - \bar{\pi})^n. \quad (6)$$

The optimal target n^* for node i is the highest integer satisfying

$$\frac{\delta}{(1 - \phi)\bar{\pi}} \left(\frac{\delta}{1 - \bar{\pi}} \right)^{n^*} \geq \frac{L}{b}. \quad (7)$$

We will choose parameters to ensure that $\delta < 1 - \bar{\pi}$, so that the LHS of (7) is decreasing in n . This is necessary for an interior solution where n^* is greater than 0 and smaller than $N - 1$.

The above describes the model under no intervention. Under an intervention, an outbreak is detected and contained with probability λ . Hence, $\Delta_n = (1 - \lambda)(1 - \phi)\bar{\pi}(1 - \bar{\pi})^n$. As a result, the optimal target n^* is the highest integer satisfying

$$\frac{\delta}{(1 - \lambda)(1 - \phi)\bar{\pi}} \left(\frac{\delta}{1 - \bar{\pi}} \right)^{n^*} \geq \frac{L}{b}. \quad (8)$$

Once the agents' optimal target is determined, the active network is formed as follows. So long as there are at least two nodes with degree less than n^* who have an available link, a link between any two nodes with degree below n^* is randomly selected to be activated. The process then stops when: 1) all links have degree n^* , or 2) all neighbours of any link with degree less than n^* have degree n^* .

A.2. Network construction and outbreak simulation

In this section, we outline the procedure for network construction and the simulation of an outbreak.

As described above, we employ the data presented in Kissler *et al.* (2020) as the basis for construction of the network. These data contain regular observations of the distance, in meters, between a large number of individuals over a period of 3 days in the town of Haslemere, UK. We specify the set of potential interactions by identifying all links where the distance between the two individuals is less than some specified distance. In the simulations conducted in this paper, the distance is either 2 or 4 m. All connections greater than this distance are removed, leaving a network of potential links. We do not factor time into this process but instead consider all interactions over the 3-day period to form the network.

The set of active links is determined through stochastic simulation. The simulation commences with no active links. Links are then selected in random order from the set of all potential links as specified by the data above. The current number of active links for each of the end points is checked. If, in both cases, the number is less than the target number of links, the link is added to the active set. Otherwise, it remains inactive. This process continues until all nodes have the target number of links or all edges have been considered. The construction of the network through this approach is stochastic and dependent on the order in which links are added. Large numbers of repetitions of the network construction process are therefore carried out and the results are calculated through a Monte Carlo simulation.

Once the network of active links is identified, outbreaks are simulated. Each outbreak is modelled iteratively. We first define two lists: L_t^C the list of those nodes currently infectious and L_t^A the list of those nodes already infected. To start the infection at step 0, a single, nonimmune node is chosen at random and added to L_0^C .

Under the no intervention case, for each node in the infectious set at period t , each of their neighbours is added to the infectious set at period $t + 1$, L_{t+1}^C , if the neighbour is not immune and is not a member of L_t^A or L_t^C . The set of already infected nodes at $t + 1$ is equal to those previously infected and those currently infectious, that is, $L_{t+1}^A = L_t^C \cup L_t^A$. As such, the infection iteratively spreads through the network until $L_t^C = \emptyset$.

Under a fast intervention, at the start of period t , each member of L_t^C is tested, and with probability λ , they are moved to L_t^A immediately, that is, prior to their neighbours being potentially infected.

Under a slow intervention, during period t , when each node is considered with probability λ its neighbours are added to L_{t+1}^A rather than L_{t+1}^C , that is, they are infected but do not infect others.

This process is repeated with different start nodes in order to simulate the effect of outbreaks and the impact of interventions on utility.

A.3. Parametric choices

Here, we give details of how the parameter values were chosen. Note that the model is constructed around a representative agent. Hence, we try to derive values of the selected parameters for the average person, while acknowledging that in reality there may be substantial heterogeneity for some of the parameters (e.g., expected losses from infection, benefits from social activity and beliefs). A similar issue emerges in relation to time. The model is static, but some of the parameters are likely to be time-varying (risk perceptions and rate of immune individuals). In this case, we consider multiple values for the same parameter to see how behaviour is likely to change during different phases of the epidemic.

- $\bar{\pi}$ (perceived probability to be infected by a new link). This parameter is related to the prevalence of the disease in the population, which changes across time depending on the phase of the epidemic. We take various measures for the prevalence of the disease. On average, between November 2020 and February 2021, 1.24 per cent of the English population was with COVID-19 every week (source: own calculations based on data from the Office of National Statistics [ONS]). This peaked in the week starting on 30 December 2020 at 2.06 per cent of the population. That said, in hotspots affected by major outbreaks, the perceived risk was probably much higher. Hence, we consider several values for $\bar{\pi}$: 0.0124, 0.0206, and multiples (0.0412, 0.0824, 0.1648).
- L (expected loss for infected individuals). Given the obvious problems in quantifying disutility from death or chronic illness/permanent injury, we restrict attention to earning losses. Hence, our measure is probably an underestimate. Average weekly pay in Britain in December 2020 was £571 (source: ONS). A worker isolating for eight working days (10 days in total) stands to lose $8 \times 571/5 = \text{£}914$. Longer COVID-19 effects are taken into account as follows. Around 1 in 5 COVID-19 patients experience symptoms after 5 weeks and around 1/10 after 12 weeks (source: ONS). We assume that half of the 1/10 who have symptoms after 12 weeks were admitted to hospital and count that as equivalent in terms of utility to losing 6 months of earnings. Hence,

$$L = 0.8 \times 914 + 0.1 \times 571 \times 5 + 0.05 \times 571 \times 12 + 0.05 \times 571 \times 24 \approx \text{£}2045. \quad (9)$$

This, however, overestimates lost earnings, since some people are able to work from home and we are not considering sick pay. If one includes statutory sick pay of approximately £96 (statutory sick pay is £95.85 in the UK) per week, lost weekly earnings are £475. Hence, $L \approx \text{£}1700$. If one further assumes that half of those self-isolating without longer term consequences are able to work from home and face no losses, we obtain $L \approx \text{£}1400$. We thus consider all three values derived above (£1400, £1700 and £2045).

- b and δ (benefit from active links and rate of decay of benefit). We use the fines imposed by the government for breaking COVID-19 rules to measure b and δ . Although these should arguably be large enough to deter most people from meeting face to face, there is some evidence of low adherence (see Smith *et al.*, 2020). This implies that, while the numbers we derive may provide an upper bound on the benefit of an additional link, the upper bound is probably tight for a significant share of the population. The fine for breaking self-isolation starts at £1000.¹³ Since September 2020, the government has restricted the maximum number of people who can meet face to face at six. The

¹³It can reach £ 10,000 for repeat offences or egregious violations.

¹⁴Some caution is needed here, since a violation of the rule of six occurs when meeting more than six people *at the same time*.

fine for breaking the rule of six is £200. We use the self-isolation fine to determine the value of the first active link ($\beta\delta$) and the rule of six fine to determine the value of the seventh active link ($b\delta^7$).¹⁴ Assuming a detection rate of 24 per cent (which, according to the Home Office, is the average detection rate across all crimes and violations), one can look at the value of the first link as $0.24 \times \text{£}1000$ and the value of the seventh link as $0.24 \times \text{£}200$. This would imply $b\delta = 0.24 \times \text{£}1000$ and $b\delta^7 = 0.24 \times \text{£}200$. Solving for b and δ , one obtains $\delta \approx 0.765$ and $b \approx \text{£}314$. For robustness, we also consider other values for δ . Note that b only enters the model via the b/L ratio and, as mentioned above, we already consider multiple values for L .

- ϕ (share of the population immune to the virus). This parameter is likely to change during the course of an epidemic as more people develop antibodies (or lose them) over time. Models using data from the Diamond Princess and accounting for asymptomatic cases estimate that approximately 2/3 of crew and passengers on the ship were not infected (Emery *et al.*, 2020). This is probably an overestimate for ϕ , since, even on a cruise ship, not everybody is necessarily exposed. More to the point, for $\phi > 0.5$, we should expect herd immunity to play a significant role. We accordingly consider different levels of ϕ between 0 and 0.5, with 10 per cent increments.
- θ . This is the probability that an outbreak hits the network (θN) divided by the number of nodes. In this version of the model with unsophisticated agents, θ may only affect behaviour indirectly, via the expectations parameter $\bar{\pi}$. Since all results presented here are conditional on an outbreak occurring and we consider different values for $\bar{\pi}$, θ can be omitted.
- λ (intervention accuracy). As this is the main parameter of interest, we let it vary between 0 and 1 with 10 per cent increments. All our results are presented for all values of λ .