

ARTICLE

GROUP TESTING AND SOCIAL DISTANCING

Spyros Galanis^{1,2*} 

¹Department of Economics and Finance, Durham University Business School, Durham, United Kingdom

²Cyprus International Institute of Management, Nicosia, Cyprus

*Corresponding author. Email: spyros.galanis@durham.ac.uk

An often overlooked strategy for fighting the COVID-19 pandemic is group testing. Its main advantage is that it can scale, enabling the regular testing of the whole population. We argue that another advantage is that it can induce social distancing. Using a simple model, we show that if a group tests positive and its members are in close social proximity, then they will rationally choose not to meet. The driving force is the uncertainty about who has the virus and the fact that the group cares about its collective welfare. We therefore propose identifying socially connected groups, such as colleagues, friends and neighbours, and testing them regularly.

Keywords: COVID-19; social proximity; health; infections; social distancing.

JEL codes: D82; D83; I12.

1. Introduction

The SARS-CoV-2 virus, which causes the COVID-19 disease, has been spreading rapidly and globally since December 2019, resulting in alarmingly high fatality rates, national lockdowns and unprecedented projected economic damage (Gopinath, 2020). The main strategies for addressing the pandemic have been individual tests, vaccinations and contact tracing. However, none of these measures has proved to be wholly effective. Countries have not managed to scale individual testing, contact tracing only works when the cases are very low, and it may take years before the majority of the world population is inoculated.

An alternative strategy, that has been proposed by several papers, is group testing (Aldridge *et al.*, 2019). A group of people is tested using a single test. If the test is negative, then no-one has the virus, whereas if it is positive, then at least one person is infected, even though their identity is not revealed. Because the samples of all individuals are mixed together, the test is anonymised, so it is impossible to know who has the virus. The main advantage of group testing is its scalability, as it enables the testing of a big part of the population on a regular basis.

In this paper, we argue that, as long as groups are selected carefully, another advantage of group testing is that it can induce social distancing. In particular, we show that if members of a group do not want to meet when all the information they have is that at least one of them is infectious, then they would still not want to meet, even if each had some private information about who is infected. This is important because, irrespective of what members of the group know about who is infectious, the public announcement of a positive group test will induce everyone not to meet. However, this does not apply to all groups, only to those that are in social proximity. We say that a group is in social proximity if its members do not want to meet when at least one of them is infectious, because their total welfare will decrease. Our results therefore show that announcing a positive test to a group in social proximity will induce members not to meet.

To provide a simple example, suppose that a group consists of Alex and Bob. If they decide not to meet, then we normalise the utility that each gets to zero. If Alex is susceptible to infection while Bob is infected, then meeting implies that Bob might be infected. Will their total utility be above 0? If it is, then it means that the cost to Alex from possibly being infected is less than the benefit to Bob from meeting. This could happen because Bob does not care about Alex's welfare. Or, it could be because Bob and Alex are complete strangers, so 'meeting' just means that they pass by each other on the street, 2 m apart, which means that Alex will probably not be infected. In that case, meeting just means that they are allowed to walk on the street, which gives both a positive utility. In that case, we say that the group is not in social proximity.

On the other hand, if meeting implies that the sum of their utilities is negative, we say that the group is in social proximity. This could be because Bob incurs an extra cost if Alex becomes infected. Or, it could be that 'meeting' means that they access the gym at the same time, so the probability that Alex is infected is big and therefore his cost is higher than the benefit of Bob from exercising.

Social proximity is not true for groups consisting of people who are complete strangers, because they do not care about each other's welfare, they meet rarely, or in a socially distanced way. For example, a randomly selected group of people in a big city is not in social proximity, either because they do not care about each other's welfare, or because meeting only takes place in a socially distanced way, if at all. This means that the sum of their utilities if they meet is greater than if they do not meet, even if someone is infected. However, the parents of students in a school could be in social proximity, as they meet regularly, either directly or through their kids. The same can be true of employees who work on the same floor of a building. As the group increases in size, social proximity reduces. For example, if the group is the whole country, the total benefit from meeting with each other will surely be greater than the cost that at least one person gets infected. On the other hand, it is important to note that social proximity is different from altruism. Members can be in social proximity without necessarily caring about the condition of other members.

We use the static version of the standard SIR epidemiology model to explain the mechanism (Anderson and May, 1992; Kermack and McKendrick, 1927). Agents can have one of the three conditions. They can be infected (I), which means that they can infect others if they meet them. They can be susceptible (S) to infection, so that they will contract the virus if they meet someone who is infected. Or, they can be recovered (R), which means that they do not transmit the virus and cannot get infected. We extend the SIR model by assuming that agents face uncertainty about what is their condition, and the condition of others. This is relevant in the current COVID-19 pandemic, as several people are asymptomatic after contracting the virus, so they may not know that they are infected (Mizumoto *et al.*, 2020). Agents have a common prior and receive private information about everyone's condition. Each agent has two available actions: to meet within the group or not to meet. Their utility depends on their condition, their action and the actions of everyone else. They rationally choose the action that maximises their expected utility, given their updated beliefs and actions of others. We examine whether the group will decide to meet in two settings: a non-strategic and a strategic one.¹

When a group test turns out positive, it is publicly announced to the group. We then say that the group is in danger, because at least one member is infected, so if they meet then it is likely that at least one member is going to be infected. A group test is anonymised and there is no way of knowing who is infected. Although the identity of the infected agent is not revealed, it becomes common knowledge that if the group meets, then the group will be in danger. Do agents care about this possibility? If they do, we say that they are in close social proximity.

In a non-strategic setting, social proximity specifies that the agents' total welfare is greater if they do not meet, given that the group is in danger, as compared to if they meet. In a strategic setting, social proximity specifies that it is *ex ante* Pareto optimal not to meet, given that the group is in danger. The *ex ante* stage is before receiving their private information, but after they are notified that the group is in

¹We also note that our model is more general than the SIR model, in the sense that we formulate our results in terms of states of the world, which could describe fewer (e.g. just S and I) or more conditions for each agent.

danger. We formalise this condition by assuming that there is no strategy profile that ex ante Pareto dominates the strategy profile where no one meets.

Will agents decide to meet? In a non-strategic setting, we show that it cannot be common knowledge that they decide to meet (Theorem 1). In a strategic setting, we show that not meeting is the unique Bayesian Nash equilibrium (Theorem 2).

To gain some intuition, consider first the following ‘dynamic’ story. Suppose that after a positive group test, a member of the group expresses eagerness to meet. Then, all others deduce that she must consider it very likely that she is not in danger, either because she currently has the virus, or because she had it in the past and is now immune. Since at least one member must have the virus, everyone else updates upwards the probability that they themselves are in danger, so this makes them even less willing to meet. If, given this updating, others are still willing to meet, the remaining members become even more cautious. Social proximity implies that not everyone can be better off from meeting, given that the group is in danger, hence as the updating of beliefs continues, eventually some members decide not to meet.

We now provide a rough sketch of the proofs of the two results. Consider first the non-strategic environment and normalise the utility of any agent from not meeting to be 0. An agent chooses to meet if, according to her private information and updated beliefs, her expected utility from doing so is strictly positive. A public announcement that meeting is dangerous for the group implies that it becomes common knowledge that at least one agent would be infected if they met. Even though agents may not know who the newly infected member might be, social proximity implies that it is common knowledge that the sum of their utilities will always be weakly negative. Can it be common knowledge that they decide to meet? If that is the case, then it is also common knowledge that everyone’s expected utility, according to their own private information, is strictly positive. But now there is an incompatibility, because while it is common knowledge that everyone always has strictly positive expected utility, it is also common knowledge that the sum of the utilities for all agents is always weakly negative. As both statements cannot be true simultaneously, the group does not meet.

In a strategic setting, suppose that there is a Bayesian Nash equilibrium where some agents decide to meet, even though it is common knowledge that the group is in danger. Take any agent and consider her individual decision problem, given that everyone else’s actions are fixed. If her strategy is never to meet, she can guarantee an ex ante expected utility of 0. Since at each state she receives some private information and not meeting is always an option, her equilibrium strategy cannot result in her getting an ex ante expected utility lower than 0, otherwise the value of information would be negative. Hence, her ex ante expected utility from her equilibrium strategy cannot be lower than 0. As this is true for all agents, ex ante Pareto optimality of not meeting, due to social proximity, implies that the equilibrium strategy is that no one meets.

The value of information result states that in a single-agent decision problem, an agent will always be better off ex ante, if in the interim stage she receives more information. Hence, the value of information is positive. The intuition is that because actions are conditioned on information, more information means that she can adjust better her actions, depending on what the true state is, so that in expectation her utility increases. This result requires that agents are sophisticated, so that their information structure forms a partition of the state space. See Galanis (2015, 2016, 2018) and Geanakoplos (1989), among others, who show that the value of information may be negative when agents are boundedly rational.

1.1. Related literature

There is growing literature on the economics of pandemics, incorporating individual decision making and notions of equilibrium. Toxvaerd (2019, 2020) extends the SIS and SIR models by introducing endogenous social distancing. Makris (2020) allows for fatalities and risk heterogeneity, calibrating the model to UK data in order to examine various government interventions. Eichenbaum *et al.* (2020)

develop an SIR-macro model to study containment policies, whereas Alvarez *et al.* (2020) examine the optimal lockdown policy.

These models are dynamic and aim to trace the spreading of the virus and the economic consequences. Our model is static and only concerns a single decision of whether to meet within a group. Moreover, an aspect that is missing from these models, that we add here, is that agents do not know what their condition is when making a decision about who to meet. This is especially relevant in the current COVID-19 pandemic, as it seems that several people are asymptomatic for several days after contracting the virus, so they do not know that they are infected. Moreover, our emphasis is more local, as we identify conditions on groups which imply that the members will not meet.

Several papers (e.g. Ferretti *et al.*, 2020; Salathé *et al.*, 2010; Yoneki and Crowcroft, 2014) study digital contact tracing, examining whether information about geographical proximity between infected and susceptible can be used, in order to decrease the spread of a virus. However, these papers do not take into account the behavioural aspect of transmission and do not try to predict how individuals will behave. The current paper incorporates the behavioural aspects of choosing to meet and argues in favour of leveraging social proximity, as an added tool in fighting pandemics.

Group testing is not a new idea and goes back to Dorfman (1943). It has been used to detect syphilis, hepatitis B and HIV, among others. See Aldridge *et al.* (2019) for a survey. Group testing can also be performed in the case of Covid-19.² Gollier and Gossner (2020) show how group testing can be optimised to multiply the efficiency of tests against Covid-19, using three applications.

Our results are closely related to the *no trade theorems*, first discussed by Aumann (1976) and Milgrom and Stokey (1982). Aumann (1976) showed that ‘we cannot agree to disagree’, whereas we show that ‘if we agree that the group is in danger, we cannot agree to meet’. Our result that the unique Bayesian Nash equilibrium is that the group does not meet, is closely related to a no trade result of Geanakoplos (1989). Several papers examine the question of whether to initiate contact between two individuals when one might be infected (Matthies and Toxvaerd, 2016; Toxvaerd, 2019, 2021). However, to the best of our knowledge, the current paper is the first that uses the ‘agreeing to disagree’ type of results in the economics of epidemiology.

The paper is organised as follows. Section 2 presents the model, formalising the notion of social proximity. Section 3 presents the two results, that it cannot be common knowledge that the group meets and that the unique Bayesian Nash equilibrium is that they do not meet. Section 4 concludes and discusses the policy implications. The proofs and the technical details are contained in the Appendix.

2. Model

Let H be a finite set of n humans, or agents. An agent can have one of three conditions: be susceptible (S), infected (I) or recovered (R). Agents are uncertain about their condition and the condition of others. Their uncertainty is summarised by state space Ω . A state of the world $\omega \in \Omega$ specifies the condition of each agent (S, I or R).³ Although agents are uncertain about the true state, they have some private information. Agent i 's private information is represented by a partition Π_i of the state space Ω . If $\omega \in \Omega$ is the true state, agent i is informed that some state in $\Pi_i(\omega) \subseteq \Omega$ is true. Agents share a common probability distribution p over Ω , so that the ex ante probability of state ω is $p(\omega) > 0$.

To provide an example, suppose there are two agents. A state ω describes the condition of both individuals, so state SI specifies that agent 1 is susceptible and agent 2 is infected. The state space Ω is $\{SS, SR, SI, RS, RR, RI, IS, IR, II\}$. Suppose that agent 1 always knows her own condition, because she does individual tests, but she has no information about the condition of agent 2. Her partition Π_1 is

²Technion—Israel Institute of Technology have shown they can test more than 60 patients simultaneously. For more details, see <https://technionuk.org/video/pooling-method-for-accelerated-testing-of-covid-19-from-technion/>.

³In general, we do not explicitly use the three conditions in any of the results, as we state them only in terms of states of the world.

$\{\{(SS,SR,SI),\{RS,RR,RI\},\{IS,IR,II\}\}$, with three partition elements. Suppose that agent 2 only knows whether she is infected or not, because of the symptoms she develops. Her partition Π_2 is $\{\{(SI,RI,II),\{SS,SR,RS,RR,IS,IR\}\}$.

Each agent i has two choices, $C = \{0,1\}$: either to meet with the people in the group, 1, or not, 0. Let $f_i : \Pi_i \rightarrow C$ be the strategy of agent i , mapping each of her partition elements to an action in C . Let $f = (f_1, \dots, f_H)$ be a profile of strategies and $f(\omega) = (f_1(\omega), \dots, f_H(\omega))$ be the particular realisation of actions at state ω .

Agent i 's utility depends on the actions and the condition (S, I or R) of everyone in the group. For example, $u_i(\omega, f(\omega))$ is i 's utility when the true state is ω and the actions of everyone is given by the strategy profile $f(\omega)$. Our first assumption normalises i 's utility to be 0 from not meeting with the group, irrespective of what everyone else is doing.

Assumption 1. For all agents $i \in H$, states $\omega \in \Omega$ and strategy profiles f with $f_i(\omega) = 0$, $u_i(\omega, f(\omega)) = 0$.

2.1. Social proximity

We now formalise the notion of social proximity. Agent i 's utility from meeting with the group depends on the condition (S, I and R) and the action of each member of the group. What does it mean that she cares about other members of the group? To provide some intuition, consider first the simple case where the group consists of two agents, i and j . If both decide to meet, agent i will get a positive payoff m_{ij} . However, if i is susceptible (S) and j is infected (I), agent i will incur a cost from contracting the virus, so her total utility will be less than if they did not meet. From the normalisation of Assumption 1, her utility will be negative. If i is infected and j is susceptible, then j will incur a cost from contracting the virus and have a negative utility. In states (SS, SR, RR, RS, RI and IR), the virus is not transmitted and both agents have positive utility. We argue that II should be in D at the end of this section.

Let $D = \{SI, IS, II\} \subseteq \Omega$ be the event that ‘the group is in danger’, because at least one member of the group will contract the virus by meeting. If the group consists of more than two agents, D is the set of all states ω such that at least one i is susceptible and at least one j is infected. For example, if there are three agents, $D = \{SIS, SII, SIR, ISS, ISI, ISR, IIS, RIS, SSI, RSI, III\}$. Recall that a group test is anonymised, so it is impossible to know who in particular is infected.

Let $\mathbf{1}$ be the strategy profile specifying that all agents decide to meet always, whereas $\mathbf{0}$ is the strategy profile where all agents decide never to meet. The following assumption specifies that for each state in D , which describes that the group is in danger, the total welfare of the group is higher if they all decide not to meet, as compared to deciding to meet.

Assumption 2. For all states $\omega \in D$, $\sum_{i \in H} u_i(\omega, \mathbf{1}) \leq 0 = \sum_{i \in H} u_i(\omega, \mathbf{0})$.

To understand this assumption, suppose first that the group consists of two agents, i, j , and $D = \{IS, SI, II\}$. At state $\omega = IS$, agent j is susceptible and therefore her utility from meeting within this group is negative, $u_j(IS, \mathbf{1}) = -k < 0$, because she may contract the virus. What about i 's utility at $\omega = IS$? It is reasonable to assume that if the two agents are friends or care about each other, i 's utility, $u_i(IS, \mathbf{1}) = l$, cannot be greater than k . In other words, if i knew with certainty that the state is IS and therefore she would surely infect j , her benefit l would not outweigh the cost k incurred by j , so that $-k + l \leq 0$ and the total welfare of the group is weakly negative.

Should state II , ‘everyone is infected’, be included in D ? If everyone is infected, is the group in danger? One could argue that it is not, because people cannot get more infected. On the other hand, it could be that there are different strains of the virus, or a member may have a high virus load that could be passed to another with a low load, so it could still be dangerous to meet. Health authorities around the world do not allow infected people to freely meet, so the latter arguments are probably more prevalent. Our results do not depend on whether II belongs to D , as Theorems 1 and 2 are stated in terms of an abstract set D . However, this matters for the interpretation of our model, because we consider that the announcement of

a positive test makes it common knowledge that the group is in danger. For the remainder of the paper, we assume that II belongs to D , or that II is impossible.

Finally, it may seem strange that within a group of strangers, where there is no social proximity, Assumption 2 implies that the total welfare increases from meeting, whereas for people that are in social proximity, the welfare decreases. There are two reasons for that. First, members of groups in social proximity may care about each other, which means that they will incur a cost if another member gets infected when they meet. Second, they meet in places where the danger of transmission is high, for example, because they work in the same building or their kids meet in the school, which means that they increase the probability that they become infected. On the other hand, strangers may not care about each other or they meet in places where the transmission is low, such as an open space park, so their total welfare increases.

3. Results

Suppose that it is announced that the test is positive, so it becomes common knowledge that the group is in danger, because at least one member might get infected if they meet. What will they do? We examine this question in two settings. In a non-strategic setting, we show that it cannot be common knowledge that they meet. In a strategic setting, not meeting is the unique Bayesian Nash equilibrium.

3.1. We cannot agree to meet

Let M be the event describing that everyone in the group decides to meet. A state ω belongs to M if each agent i 's expected utility is strictly positive, given that everyone else also decides to meet.⁴ Because the utility from not meeting is 0, a strictly positive expected utility implies that the individual would like to meet. If the group decides to meet at $\omega \in M$, this becomes common knowledge. That is, we assume that members do not meet privately with each other, but everyone meets with everyone else and this is common knowledge. Formally, we say that the group agrees to meet at state $\omega \in \Omega$ if M is common knowledge at ω .⁵ This implies that if only a few members meet, we do not consider that the group has met.

This definition can also be interpreted as another aspect of social proximity. If the group is very large and members do not know each other, or they do not have a common place where they meet, it becomes more difficult to monitor what everyone is doing, hence even if they meet, this does not become common knowledge, so they cannot agree to meet.

Suppose that the group test turns out positive and this is announced within the group. This means that the event D , 'the group is in danger', becomes common knowledge. More generally, in all states in D , the group test turns out positive and event D is common knowledge. Will the members of the group decide to meet? As the following Theorem shows, the answer is no.

Theorem 1. *Under Assumptions 1 and 2, if at all states in D it is common knowledge that the group is in danger, the group cannot agree to meet at any state in D .*

3.2. The unique equilibrium is not to meet

We now examine the same question in a strategic setting. Suppose that after the group is publicly notified that they are in danger, so D is common knowledge, they play the following standard Bayesian game. A state $\omega \in D$ occurs, each agent receives her private information, updates her beliefs and plays a best response. Could it be that meeting is a Bayesian Nash equilibrium? We show that the unique equilibrium is that they do not meet.

⁴Recall from Assumption 1 that utility from not meeting is 0, irrespective of what everyone else is doing.

⁵See the Appendix on how we define common knowledge.

Social proximity, defined as the property that members care about their collective welfare, is expressed in this game by specifying that it is ex ante Pareto optimal not to meet, when the group is in danger. The ex ante refers to the stage where the agents have not yet received their private information, but they are informed that the group is in danger, so that it is common knowledge that one state in D is true. In other words, there is a veil of ignorance about who might be infected if the group meets.

Recall that a strategy f_i of agent i maps each partition cell to one of two actions: meet (1) or not meet (0). Given a profile f of strategies, agent i 's ex ante expected utility is $\sum_{\omega \in D} p(\omega)u_i(\omega, f(\omega))$. If

$\sum_{\omega \in D} p(\omega)u_i(\omega, f(\omega)) \geq 0$, then agent i can do weakly better with f , than if she decides not to meet at

all states $\omega \in D$.⁶ Ex ante Pareto optimality implies that although this could be true for some agents, it cannot be true for all, because then it would be collectively better for the group to meet, even though they are in danger. Equivalently, ex ante Pareto optimality implies that if everyone's ex ante expected utility (given D) is weakly greater than 0, then it is exactly zero and everyone chooses not to meet.⁷ We formalise this in the following Assumption.

Assumption 3. *If $\sum_{\omega \in D} p(\omega)u_i(\omega, f(\omega)) \geq 0$ for all $i \in H$, then $f_i(\omega) = 0$ for all $i \in H$ and all $\omega \in D$.*

The following Theorem shows that if it is common knowledge that the group is in danger, then the unique Bayesian Nash equilibrium is not to meet.

Theorem 2. *Under Assumptions 1 and 3, if at all states in D it is common knowledge that the group is in danger, then the unique Bayesian Nash equilibrium is not to meet.*

4. Discussion and policy implications

Our results show that group testing can leverage social proximity and incentives, in order to induce social distancing. If a group cares about their collective welfare, they will rationally choose not to meet, as soon as they learn that some members are in danger of contracting the virus. This is true both in a strategic and a non-strategic setting. If we also consider that group testing can scale considerably faster than individual testing, it is evident that it can act as a complementary strategy for addressing the pandemic, alongside the existing ones, such as individual testing, contact tracing, lockdowns and inoculations. This is especially relevant for countries which cannot inoculate a big part of their population fast enough.

The policy implications are straightforward. As a first, step, we propose identifying groups of people that are socially and geographically connected, such as students within a school, colleagues and co-workers within a workplace, or neighbours, and testing them regularly. The geographical proximity of the members of the group makes regular testing easier to implement. More importantly, geographical proximity can induce social proximity, as it forces the group to care more about its total welfare, because meeting when at least one member is infected may impact everyone.

On the other hand, group testing has some limitations. Although it can induce social distancing within a group, it may not be as effective in limiting contact across different groups. For example, although a worker may choose not to meet his colleagues if a positive group test is announced in the workplace, he may still be willing to play football with his friends, as this group is not tested or has tested negative. This creates a trade-off, between the scalability of group testing and its possibly reduced effectiveness across groups, which would be an interesting direction for future research. It also raises the question of what is the optimal way of choosing which further groups to test, based on the currently positive tests and the membership of people across groups.

We conclude by making some comments on the interpretation of the model. First, although the SIR model is a good starting point for thinking about how the condition of agent i might impact agent j if they

⁶Recall, from Assumption 1, that she gets 0 if she does not meet, irrespective of what the others are doing.

⁷We are implicitly assuming that there is no profile of strategies, other than $\mathbf{0}$, that gives a utility of 0 to everyone.

meet, we do not use the three conditions explicitly in our model. Instead, we formalise our assumptions and results within an abstract state space, so we can accommodate fewer or more conditions. Second, we do not take a stance on whether a susceptible to infection individual will surely get infected if the group meets. What matters for our results is how their utility will decrease, if they meet when the group is in danger.

Second, the model does not preclude that some members develop symptoms. When some members have symptoms, they may know that they are infected and maybe some other members know that too. This would mean that the group is in danger, even without a positive test. In general, any private information is allowed, so members can know something about the condition of others or of themselves. However, since the model is static, it cannot describe a dynamic process where some members become symptomatic and then they get quarantined. In such an extension, the quarantined members would be removed from the group for some periods. If there is correlation between someone developing symptoms and another one becoming infected (e.g. a spouse), then the quarantine could provide public information about the condition of members that are not currently quarantined.

Finally, we discuss our implicit assumption that the announcement of a positive group test implies that D becomes common knowledge. If we consider all possible combinations of the SIR conditions for all agents, then there are two cases where a positive group test does not necessarily imply that someone will get infected. The first is that everyone is infected. As we argue in Section 2.1, different agents may have different variants of the virus, or different virus loads, so meeting can still be dangerous. The second is that one member is infected but everyone else is recovered, hence they are immune to being infected. We implicitly assume that this is not possible. One could justify this assumption by saying that even if everyone else is vaccinated, there is still the possibility that someone will be infected, so that the probability that everyone is immune is zero. Alternatively, we could say that if everyone else is 100 per cent immune, they will know it so there is no uncertainty about who is infected, and the model is trivial.

Acknowledgements. I thank Anastasios Karantounias, Miltos Makris, Flavio Toxvaerd, the editor, two referees and participants at the National Institute of Economic and Social Research (NIESR) 2021 workshop for helpful comments.

References

- Aldridge, M., Johnson, O. and Scarlett, J. (2019), 'Group testing: An information theory perspective', arXiv Preprint, arXiv:1902.06002.
- Alvarez, F., Argente, D. and Lippi, F. (2020), 'A simple planning problem for covid-19 lockdown', Working paper 2020-34, Becker Friedman Institute for Economics, University of Chicago, Chicago, Illinois.
- Anderson, R.M. and May, R.M. (1992), *Infectious Diseases of Humans: Dynamics and Control*, Oxford: Oxford University Press.
- Aumann, R. (1976), 'Agreeing to disagree', *Annals of Statistics*, **4**, pp. 1236–9.
- Dorfman, R. (1943), 'The detection of defective members of large populations', *The Annals of Mathematical Statistics*, **14**, 4, pp. 436–40.
- Eichenbaum, M.S., Rebelo, S. and Trabandt, M. (2020), 'The macroeconomics of epidemics', Technical report, National Bureau of Economic Research, Cambridge, Massachusetts.
- Ferretti, L., Wymant, C., Kendall, M., Zhao, L., Nurtay, A., Abeler-Dörner, L., Parker, M., Bonsall, D. and Fraser, C. (2020), 'Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing', *Science*, **368**, 6491, eabb6936.
- Galanis, S. (2015), 'The value of information under unawareness', *Journal of Economic Theory*, **157**, pp. 384–96.
- Galanis, S. (2016), 'The value of information in risk-sharing environments with unawareness', *Games and Economic Behavior*, **97**, pp. 1–18.
- Galanis, S. (2018), 'Speculation under unawareness', *Games and Economic Behavior*, **109**, 598–615.
- Geanakoplos, J. (1989), 'Game theory without partitions, and applications to speculation and consensus', *Cowles Foundation discussion paper*, no. 914.
- Geanakoplos, J. (1992), 'Common knowledge', *The Journal of Economic Perspectives*, **6**, 4, 53–82.
- Gollier, C. and Gossner, O. (2020), 'Group testing against covid-19', Technical report, EconPol Policy Brief.
- Gopinath, G. (2020), 'Limiting the economic fallout of the coronavirus with large targeted policies', available online at <https://blogs.imf.org/2020/03/09/limiting-the-economic-fallout-of-the-coronavirus-with-large-targeted-policies/>

Kermack, W.O. and McKendrick, A.G. (1927), ‘A contribution to the mathematical theory of epidemics’, *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, **115**, 772, pp. 700–21.

Makris, M. (2020), ‘Covid and social distancing with a heterogeneous population’, *Economic Theory*, forthcoming.

Matthies, K. and Toxvaerd, F. (2016), ‘Rather doomed than uncertain: Risk attitudes and sexual behavior under asymptomatic infection’, Working paper.

Milgrom, P. and Stokey, N. (1982), ‘Information, trade, and common knowledge’, *Journal of Economic Theory*, **26**, pp. 17–27.

Mizumoto, K., Kagaya, K., Zarebski, A. and Chowell, G. (2020), ‘Estimating the asymptomatic proportion of coronavirus disease 2019 (covid-19) cases on board the diamond princess cruise ship, Yokohama, Japan, 2020’, *Eurosurveillance*, **25**, 10, p. 2000180.

Salathé, M., Kazandjieva, M., Lee, J.W., Levis, P., Feldman, M.W. and Jones, J.H. (2010), ‘A high-resolution human contact network for infectious disease transmission’, *Proceedings of the National Academy of Sciences*, **107**, 51, pp. 22020–5.

Toxvaerd, F. (2019), ‘Rational disinhibition and externalities in prevention’, *International Economic Review*, **60**, 4, pp. 1737–55.

Toxvaerd, F. (2020), ‘Equilibrium social distancing’, Cambridge-INET working paper series no: 2020/08.

Toxvaerd, F. (2021), ‘Contacts, altruism and competing externalities’, Working paper.

Yoneki, E. and Crowcroft, J. (2014), ‘Epimap: Towards quantifying contact networks for understanding epidemiology in developing countries’, *Ad Hoc Networks*, **13**, pp. 83–93.

A. Appendix

A.1. Preliminaries

The state space Ω is a subset of the Cartesian product $\{S, I, R\}^H$. At state $\omega \in \Omega$, agent i updates her beliefs using Bayes’ rule, so that she assigns probability $\frac{p(\omega')}{p(\Pi_i(\omega))}$ to ω' if $\omega' \in \Pi_i(\omega)$ and 0 otherwise. Let $u_i : \Omega \times C^H \rightarrow \mathbb{R}$ be i ’s utility, a function of the state and everyone’s action. At state ω and given a profile of strategies f , agent i updates her beliefs using her private information $\Pi_i(\omega)$ and decides to meet with the group if her expected utility is strictly greater than if she does not. Formally, i decides to meet at ω if $\sum_{\omega' \in \Pi_i(\omega)} \frac{p(\omega')}{p(\Pi_i(\omega))} u_i(\omega', 1, f_{-i}(\omega')) > \sum_{\omega' \in \Pi_i(\omega)} \frac{p(\omega')}{p(\Pi_i(\omega))} u_i(\omega', 0, f_{-i}(\omega'))$, where $f_{-i}(\omega')$ is the profile of actions at ω' , for all agents except i .

Let the profile of actions where no one meets at any state to be $\mathbf{0}$, so that $\mathbf{0}_i(\omega) = 0$ for all $i \in H$ and $\omega \in \Omega$. The profile of strategies where everyone meets always is denoted $\mathbf{1}$, so that $\mathbf{1}_i(\omega) = 1$ for all $i \in H$ and $\omega \in \Omega$. To simplify the notation, we write i ’s utility as $u_i(\omega, \mathbf{1})$, instead of $u_i(\omega, \mathbf{1}(\omega))$.

An event E is a subset of Ω . For example, the event ‘agent j meets with the group’ is the set of states ω' such that $f_j(\omega') = 1$. Agent i knows event E at ω , if $\Pi_i(\omega) \subseteq E$. This means that in all states that she considers possible at ω , E is true. Let M be the event describing that everyone in the group decides to meet. A state ω belongs to M if each agent i ’s expected utility is strictly positive, given that everyone else also decides to meet. Formally, M consists of all states ω such that $\sum_{\omega' \in \Pi_i(\omega)} \frac{p(\omega')}{p(\Pi_i(\omega))} u_i(\omega', \mathbf{1}) > 0$, for all $i \in H$. The event ‘agent j knows that everyone meets within the group’ consists of all states ω such that $\Pi_j(\omega) \subseteq M$.

In order to define higher orders of reasoning about the knowledge of others, let $\Pi_i(F) = \bigcup_{\omega' \in F} \Pi_i(\omega')$ be the set of all states that i might think are possible, if the true state is in F . Using this notation, we can say that $\Pi_j(\Pi_i(\omega))$ is the set of states that, at ω , agent i thinks that j considers possible. If $\Pi_j(\Pi_i(\omega)) \subseteq F$, then we say that i knows that j knows F . An event E is common knowledge at ω if $\Pi_{i_n}(\Pi_{i_{n-1}} \dots (\Pi_{i_1}(\omega))) \subseteq E$, for any sequence of agents i_1, \dots, i_n .⁸

We restrict the definition of Bayesian Nash Equilibrium to states in D , because D is common knowledge after the public announcement of the positive test. A profile of strategies f is a Bayesian Nash equilibrium if, for all states $\omega \in D$, f_i is a best response for agent i . Formally, we have that

$$\sum_{\omega' \in \Pi_i(\omega)} \frac{p(\omega')}{p(\Pi_i(\omega))} u_i(\omega', f_i(\omega), f_{-i}(\omega')) \geq \sum_{\omega' \in \Pi_i(\omega)} \frac{p(\omega')}{p(\Pi_i(\omega))} u_i(\omega', c, f_{-i}(\omega')) \text{ for all } c \in \{0, 1\}.$$
⁹

⁸These notions are explained clearly in Geanakoplos (1992).

⁹Recall that a strategy f_i maps elements of Π_i to C , hence $f_i(\omega') = f_i(\omega'')$ for all $\omega', \omega'' \in \Pi_i(\omega)$.

A.2. Proofs

Proof of Theorem 1. Suppose that at all states in D it is common knowledge that the group is in danger, but the group decides to meet at $\omega \in D$. By definition, event M is common knowledge at ω .

Let $\mathcal{M}(\omega)$ be the set of states that are reachable from ω . Formally, $\mathcal{M}(\omega)$ is the union of sets $\Pi_{i_n}(\Pi_{i_{n-1}} \dots (\Pi_{i_1}(\omega)))$, for any sequence of agents i_1, \dots, i_n . Say that an event is self-evident (within the group) if whenever it occurs, everyone knows it. Formally, if $\omega' \in E'$, then $\Pi_i(\omega') \subseteq E'$ for all $i \in H$. Then, $\mathcal{M}(\omega)$ can be described as the smallest self-evident event that contains ω .¹⁰ Aumann (1976) shows that an event E is common knowledge at ω , if and only if $\mathcal{M}(\omega) \subseteq E$.

We therefore have that $\mathcal{M}(\omega) \subseteq M \cap D$, which implies that for each $i \in H$ and all $\omega' \in \mathcal{M}(\omega)$, $\sum_{\omega'' \in \Pi_i(\omega')} \frac{p(\omega'')}{p(\Pi_i(\omega'))} u_i(\omega'', \mathbf{1}) > 0$. Because $\mathcal{M}(\omega)$ is an element of the finest common coarsening of the partitions of all agents within the group, we have that for each $i \in H$, $\mathcal{M}(\omega)$ is partitioned by some elements of i 's partition Π_i . By noting that for every such element $\Pi_i(\omega')$, we have $\sum_{\omega'' \in \Pi_i(\omega')} \frac{p(\omega'')}{p(\Pi_i(\omega'))} u_i(\omega'', \mathbf{1}) > 0$, and adding over all these elements, we get $\sum_{\omega' \in \mathcal{M}(\omega)} p(\omega') u_i(\omega', \mathbf{1}) > 0$. By adding over all agents in H , we have that $\sum_{\omega' \in \mathcal{M}(\omega)} p(\omega') \sum_{i \in H} u_i(\omega', \mathbf{1}) > 0$. This implies that for some state $\omega' \in \mathcal{M}(\omega)$, we have that $\sum_{i \in H} u_i(\omega', \mathbf{1}) > 0$. But this contradicts Assumption 2 and the fact that $\mathcal{M}(\omega) \subseteq D$. Hence, the group cannot agree to meet at any $\omega \in D$. \square

Proof of Theorem 2. The proof is similar to that of Theorem 3 of Geanakoplos (1989).

Suppose that at all states in D it is common knowledge that the group is in danger, so one state $\omega \in D$ is true. Let (f_1, \dots, f_H) be a Bayesian Nash equilibrium. Fix f_j for all $j \neq i$ and look at the one-agent decision problem for agent i . By choosing 0 at all states in D , all types of i can guarantee a payoff of 0, from Assumption 1, irrespective of what other players are doing. Because each type of agent i plays best response given f_{-i} , we have that, for each $\omega \in D$, $\sum_{\omega' \in \Pi_i(\omega)} \frac{p(\omega')}{p(\Pi_i(\omega))} u_i(\omega', f(\omega')) \geq 0$. Adding over all partition elements of D , we have $\sum_{\omega' \in D} p(\omega') u_i(\omega', f(\omega')) \geq 0$. Because this is true for all agents, Assumption 3 implies that $f_i(\omega) = 0$ for all $i \in H$ and all $\omega \in D$, hence no-one chooses to meet.

¹⁰The partition generated by M is called the finest common coarsening of the partitions of all agents.