

Comment on “Integrated epi-econ assessment of vaccination,” by Boppart, Harmenberg, Krusell, and Olsson

Marina Azzimonti
Stony Brook University and NBER

1 Introduction

Boppart et al. (2021) present a macro SIR model with endogenous behavioral responses, in which individuals can adjust consumption and labor in sectors or activities that are perceived to be more risky (e.g. restaurants). I refer to these behavioral responses as the ‘fear factor’ effect. Because agents respond to aggregate infection risk, there is an interaction between the level of economic activity (the ECON part of the model) and the spread of infections (the EPI part of the model). Higher risk induces agents to substitute away from ‘social leisure’ (e.g. in-person leisure activities or consumption near other agents).

The authors show that the degree of heterogeneity among agents, and in particular their age, is relevant in designing optimal lockdown and vaccination policies. A key finding is that due to the presence of externalities, a decentralized vaccination strategy yields sub-optimal results. Households would prefer their older members to receive a vaccine first, because these are the ones subject to higher risk of severe illness or death. A planner, instead, would vaccinate younger members of the household before vaccinating older ones. Why? Because young agents do not internalize how their choices affect the decisions of older individuals in the population as a whole. As a result, young agents engage in too much economic activity, spreading the disease faster, whereas older individuals reduce their probability of infection by staying inside their homes. By vaccinating young individuals first, society is able to reduce the number of infections and deaths, while at the same time permitting older agents to increase their consumption of social leisure. Notice that this result would not arise in an environment that does not feature ‘fear factor’ effects.

In this note, I briefly describe their model in Section 2, and contrast some of their key assumptions with empirical evidence. In Section 4, I discuss how taking into consideration the stability of human relationships can affect their results, drawing insights from network models. Section 5 concludes.

2 Brief Description of the Model

The top panel of Figure 1 illustrates the timeline of a typical day in this model. Because sleeping time is unaffected by the disease, a day is normalized to have one unit of time, to be devoted between working activities n and leisure activities h . Working provides income, and this allows agents to buy consumption goods c , with preferences defined over those two goods.




Figure 1: Timeline: a day

Differently from a standard consumption-leisure model, there are two sectors in the economy with differential risk of infection (see bottom panel in the figure). The V sector (for private) involves the consumption of goods in isolation. An example would be a take-out frozen pizza bought at the supermarket. Production in this sector is also low risk, because the good is produced by a team of manufacturing workers who do not interact with customers. The B sector (for public) involves goods that are consumed and produced in the proximity of other agents. An example would be eating at a pizzeria or working there as a waiter. Complementarities between consumption and leisure of these goods in these two sectors imply that agents will be able to imperfectly substitute one for another when infections are high. That is captured by an aggregate nested-CES utility function

$$U^i(c, h) = CES(CES(c_B^i, h_B^i), CES(c_V^i, h_V^i)),$$

where i represents the age of an individual (e.g. $i \in \{\text{young, old}\}$). It is assumed that older agents are relatively less productive than younger ones in both sectors, a force that will induce older agents to work less when infections are high. Choices are made at the household level, with households composed by a continuum of young and old individuals.




Figure 2: Choices

Figure 2 illustrates the set of choices, and how they relate to infection risk. Working from home and consuming at home (V sector) are the safest choices. Working at a factory ‘in person’ has intermediate risk levels associated to it. The riskiest activities involve working or consuming in the B sector. In general, the level of risk that individuals of age i face depends the ‘rate of transmission,’ a function of the number and intensity of people engaging in each activity (through work and/or leisure time devoted to it)

$$\frac{T^i}{S^i} = \pi_B[n_B^i + h_{B,w}^i] + \pi_V n_{V,w}^i.$$

Hence, the level of economic (ECON) activity directly affects the evolution of the pandemic (EPI). The pandemic, in turn, affects economic activity through the endogenous responses of households (via n^i and h^i). The infection risks per unit of time in each sector, π_B and π_V are determined in equilibrium by the aggregate decisions of all households in the economy.

The authors show that in a decentralized equilibrium old agents substitute towards the V sector (in terms of consumption and hours worked) as infections rise, in order to shield themselves from risk. As a result, their degree of leisure rises. When making decisions, households ignore the effect of their private actions on π_B and π_V . This is the main source of the externality. Figure 12 in Boppart et al. (2021)’s paper illustrates the key implication of this: a planner would reduce the number of hours that the young spend on sector B while at the same time increasing the social leisure of the old. An interesting implication is that even though this results in *less* output, the flow in utility would be higher, increasing aggregate welfare relative to a decentralized equilibrium. This is the case because the planner does not only care about maximizing output, but about also maximizing welfare. And higher flow utility can be attained by letting older agents going to restaurants. In terms of vaccinations,

households would prefer vaccines to be given to old agents first, because of their higher likelihood of death. A planner, who internalizes how vaccinations affects π_B and π_V —and hence reduce aggregate risk—would prioritize vaccinating the young earlier. Households, instead, only consider individual risk. A key lesson from this paper is that policymakers should take general equilibrium effects into consideration when designing vaccination strategies.

3 Thoughts and Comments

This is a very ambitious project: the authors embedded SIR dynamics in a model in which agents react endogenously to the disease, and characterize both, the decentralized and the planning solutions to the dynamic problem. In what follows, I summarize some thoughts, comments, and potential future avenues of research for this paper.

3.1 Is the Fear Factor real?

A premise of this paper is that individuals have a high elasticity of substitution between activities in different sectors, and that they respond significantly to infection risk. The first question that comes to mind is whether this is indeed the case. The experience of New York City at the outset of the pandemic provides significant evidence to that effect.

The top panel of Figure 3 displays the timeline of lock-down policies in New York in early March 2020. The first case of COVID-19 was discovered on March 1st. A state of emergency was declared a week after, and distanced learning started on March 11th for higher education. Elementary schools and high-schools shut down on March 16th. The ban on customer-facing establishments (e.g. the B sector) was imposed around the same date. The bottom panel shows how residents of NYC changed their behavior, measured by trends in mobility data in the area (collected from GPS cellphone locations by Google). Clearly, the reduction in workplace and retail mobility started earlier than the lockdown policies, with people substituting away from risky activities and spending more time in their residences. Part of the reduction in work mobility was due to a shift towards ‘working from home,’ which was possible in several non customer-facing sectors (a channel emphasized in the paper as well). These trends can be observed in other areas in which infections start to rise but lockdown policies were not yet in effect. This suggest that fear factor effects are present and induce changes in individual behavior (both in terms of consumption and leisure) as predicted by the model.



Figure 3: Timeline (top) and Google Mobility (bottom) in NYC.

3.2 Does age matter?

Another important assumption of the model is that young and old agents react differently to infection risk. The underlying assumption is that because old agents are more likely to suffer severe complications from COVID-19 and experience higher mortality rates, they are more elastic than young individuals to increases in infection risk.




Figure 4: COVID-19 infections by age (as % of population group)

Is this a plausible assumption? In Figure 4, we can see the percentage of positive cases in two age groups—the young, aged 18-44 (in green), and the old, aged 65 and above (in red)—in New York (solid line) and Florida (dashed line).¹ The horizontal axis indicates the number of days from the date in which a 0.5% total case rate was detected in each state. The figure is displayed in this way because the pandemic hit the NY area first, and Florida much later. By the time in which infections were high enough in Florida, its citizens were already well aware of the consequences of becoming infected as they had observed what had happened in NYC. We see that the level of infections for young individuals in Florida and NY are comparable, but the level of infections among old individuals is much lower in Florida. It is possible that there are other factor affecting the gap, such as differences in the number of old individuals in nursing homes, family structure, and agglomeration. However, when considering this together with the behavior in Google mobility in Florida (not displayed, available upon request), we can conclude that the fear factor was a significant element on the different outcomes in the two states. It would be interesting to test this using county-level data, in order to quantify how much of the gap can be attributed to the fear factor.

3.3 Key simplifying assumptions

There are two key simplifying assumptions that were made in order to make the paper tractable. The first one is that there is perfect insurance within the family. The second one is that all individuals are asymptomatic. Let us consider the implications of each assumption, in turn.

Perfect insurance within the family: Consumption flows between young and old agents can be redistributed at no cost within the family. This implies that only the health benefits or reallocating effort across generations are considered in the quantitative calibration, but not the cost that would arise under imperfect insurance. For example, young agents would have a much smaller consumption level if they were unable to work (or reduce the number of hours worked) during the pandemic. Old agents would not be able to afford social leisure if they did not receive transfers from young agents in the household. The effect of market incompleteness in macroeconomic models with contagion has been studied in recent work by Hur (2021) and Glover et al. (2021). From their findings, I conjecture that the age-targeted gains associated to lock-downs and vaccination policies may be overstated in the current paper. I believe, however, that the qualitative results are likely to remain unchanged.

¹This is computed as the ratio between the number of cases in each group and the number of individuals in that group. The data was constructed by Alessandra Folgi, Fabrizio Perri, Mark Ponder and myself for an NBER presentation.

Everyone is asymptomatic: An implicit assumption in this environment is that any agent who gets sick is never sick enough to stop working or shopping. Production in the model is affected by people who *choose* not to work out of fear, but not by people who feel too sick to go to work. As a result, GDP is independent of the percentage of infected individuals. Total labor in the model is computed as

$$Labor_{j,w}^i = [S^i + I^i + R^i]n_{j,w}^i.$$

I would expect that productivity when infected is lower, or even zero, for individuals with severe effects (e.g, those who need to be treated at the hospital or stay in bed until they recover). This would imply that we would only consider $(1 - \phi)I$ infected individuals as being productive, with the parameter ϕ calibrated to the percentage of symptomatic agents in the population. A similar adjustment could be done to the level of consumption and leisure when symptomatic. These would introduce an additional feedback effect from infections to welfare and decisions, as well as aggregate variables in general equilibrium.

Another consequence of assuming that everyone is asymptomatic is that it presumes that tests do not exist and people never quarantine. As a result, individuals do not change their behavior when they get sick, which would potentially distort the EPI=>ECON channel. The advantage of this assumption is that only aggregate state variables are needed when solving the household problem, without the need to carry individual-specific states and knowing the distribution of health outcomes at each point in time. This is a very smart and useful assumption, which has major implications for the feasibility of computing the decentralized equilibrium (given the lack of ‘aggregation’), because the authors do not need to keep track of the whole distribution of states. In future research, it would be interesting to consider a situation in which agents do not know the individual states of others, but can use aggregates as proxies to the severity of infections. This would be similar to computing a Krusell-Smith model where only averages enter expectations.

Knowing your type: In this model, individuals know what type of agent they are, and how likely this type is to get sick. But they do not understand their specific likelihood of getting covid. Because only aggregate states matter for individual decisions, each person has the same probability of getting sick conditional on their type. This implies that infection risk is heterogeneous (in age, time devoted to consumption, work, and leisure in each sector) but not individual specific. For example, a waiter that has been in contact with an infected customer may choose to self-quarantine whereas another identical waiter that was not in contact would continue to work as usual. Note that the only difference assumed between these agents is their individual state (whom they have been in contact with). The ECON=>EPI channel can be distorted as a result of this assumption.

4 Stability of social interactions

As standard in SIR models, the stability of social interaction in families, leisure activities, and workplaces is not taken into consideration in this model. For example, even though young and old agents live in the same household, they do not infect each other. We know, however, that family ties are *close and stable*: individuals are in the same family for long periods of time and enjoy leisure activities together. We would expect that two old individuals living in two identical households would have different probabilities of infection if a young agent is infected in one of the families but healthy (ie susceptible) in the other. Why? Because repeated interaction within the family would make the infected individual transmit the disease with high probability in a short period of time.

The same argument can be made about workplaces: the V sector is likely to involve individuals who work in teams where the same coworkers meet every day and have limited access to customers. If at some point one of the coworkers becomes sick, that specific factory becomes a very high risk environment (especially given lack of quarantine) with most members of the team becoming sick soon after. While the paper captures that more contacts involve higher risk, it does not take into consideration that a person may be at high risk simply because one of her close contacts (family or team member) is infected.

Finally, the ‘customer base’ in the B sector is also important. Whether convenience stores have a local customer base (e.g., neighborhood clients shopping repeatedly in the same establishment) or a global customer base (e.g., a store located in an airport or shopping mall) matters for the aggregate infection risk.

4.1 Network models

Network models take these dimensions into account by considering links that vary in intensity (time spent together), stability (the frequency of interactions between any two individuals), and reach (local versus global links). In Azzimonti et al. (2021), we construct a multi-layered network model of social interaction and disease transmissibility in order to quantify the relative importance of these factors. We find that in local stores (with repeated and stable interactions within a neighborhood) there is little transmission of the disease. Large events and stores with a global customer base (where interactions between individuals from different social circles is more random), on the other hand, result in faster disease transmission. This happens even though the ‘fear factor’ is absent. Over time, the network reaches *local herd immunity*, where the disease dies out in social circles that have a low frequency of meetings with people outside their circles. This, in turn, can lead to a slow burn (where the disease spreads slowly but never fully disappears) and can give rise to *infection waves*. These arise when a few agents in different social circles attend a large event and spread the infection back to their close contacts, generating a large increase in infections over time.

4.2 A simple network model

To illustrate how the stability of connections affects the evolution of the disease, I present a simplified version of the network model in Azzimonti et al. (2021). I abstract from the fear factor, age differences, and economic activity. The objective of this section is to formalize the ideas discussed above, and compare the evolution of the disease in a network model to those in a standard SIR one. Part of the following material was developed jointly with A. Fogli, F. Perri and M. Ponder.

Time is discrete and denoted by $t = \{1, \dots, T\}$. At each point in time, there are M individuals interacting on a social network. Individuals, sometimes referred to as ‘nodes,’ are connected by ‘links,’ representing a bidirectional interaction between them. Mathematically, the set of connections are specified by a graph of size M (e.g. with M nodes) through an $M \times M$ adjacency matrix $\mathbf{G}$, with elements $g_{i,j} = 1$ if agent i is connected with agent j , and zero otherwise and $g_{i,i} = 0$ for all i . Following recent contribution in the network literature (as in Acemoglu et al. (2010) and in Azzimonti and Fernandes (2018)), I assume that not all connections are active all the time. In particular we can define the connection between nodes i and j to be active if $\alpha_{ijt} = 1$, where $\alpha_{ijt} = 1$ is an i.i.d Bernoulli random variable with mean $0 < \bar{\alpha} < 1$.

Each agent, or node of the network, has an individual health state. The health state can be, as in the SIR set-up, Susceptible, Infected or Recovered. When node i is susceptible at time t , it can become infected at time $t + 1$ if it has active connections with other infected nodes. Assuming independence of each infection event, the probability of infection of a susceptible node i at time t is given by


$$p_{i,t}(\mathbf{G}_t, \pi) = 1 - \prod_{j=1}^M (1 - \pi)^{\alpha_{ijt}g_{i,j}I(j,t)} \quad (1)$$

where π here is the probability that a susceptible node gets infected when it meets an infected node, $I(j,t)$ is the function that is 1 if node j is infected at time t and 0 otherwise. When a node is infected at time t , it recovers at time $t + 1$ with probability ρ . Recovery is an absorbing state. Note that the term $\alpha_{ijt}g_{i,j}$ captures the common network structure which determines the spread of the infection.

In order to highlight the importance of the network structure, we can consider two standard network topologies: a ‘complete network’ and a ‘small world network.’ In the complete network, each node is connected to every other node, but only a fraction $\bar{\alpha}$ of connections is active at any given point in time. An implication of this is that each node is actively connected to a random sample—of fixed size $\bar{\alpha}(M - 1)$ —of the whole population.

The complete network is illustrated in Panel (a) of Figure 5, for a case in which there are $M = 20$ nodes. The thick solid lines represent the active contacts of node i at time t (i.e. the i, j connections for which $\alpha_{ijt}g_{i,j} = 1$), while the thin lines represent inactive connections (i.e. the i, j connections for which $g_{i,j} = 1$ and $\alpha_{ijt} = 0$). The value of $\bar{\alpha}$ is chosen such that the expected number of active connections for each node, denoted as C , is equal to 3. Because each node has active connections with a random sample of the whole population, the probability of infection is proportional

(a) A complete network, $M=20$, $C=3$



(b) A small world network, $M=20$, $C=3$




Figure 5: Network structure (solid: active connections, dashed: inactive connections).

to the fraction of infected nodes in the population. As a result, this network has the same infection dynamics than a standard SIR model.

In the small world network (see Watts and Strogatz (1998) and Fogli and Veldkamp (2020) for an economic application), each node is connected to its immediate neighbors to the right and to the left, and some nodes which are distant in the network. Panel (b) of Figure 5 illustrates an example in which each node is connected to 4 nodes that are close and two nodes that are distant. The value of $\bar{\alpha}$ is chosen, as before, such that the expected number of active connections in each period is $C = 3$.

Simulating infection dynamics: In order to illustrate the evolution of the disease, I parameterize the model using: (i) an infection probability upon contact of $\pi = 12\%$, (ii) a probability of recovery $\rho = 8\%$, an initial prevalence of the disease of $s_0 = 1\%$, (iv) a network of size $M = 2000$, (v) an expected number of contacts $C = 4$ and, in the case of the small network world, I further assume that each agent has 4 short links (local contacts) and 4 long links (global contacts). While the number of active contacts is the same in the two networks, the connections in the small world network are more local and more stable. In particular, each agent draws 4 connections each period from the same pool of 8 possible links, whereas in the complete network each agent draws 4 connections to individuals all over the network. This will generate significantly different dynamics for the spread of the disease. The pandemic is assumed to hit this economy on day 10 of the simulation.

Figure 6 shows the proportion of infected individuals in the population (e.g. the case rate I/M) over time in the complete network model (labelled SIR) and in the small-world network (labeled NET) for the simulation. Because connections are more random in SIR, the infection grows faster early on than in NET, as infected nodes




Figure 6: Infection dynamics in SIR and Network

spread the disease to many susceptible nodes. As the recovered nodes become a majority, the percentage of infected rapidly declines. Contacts are more local in NET, implying that the infection develops fast only in specific pockets (social circles) at the outset, and travels more slowly through the network. Over time, the infection takes longer to decline than in the SIR model. This is the case because there are still pockets of susceptible nodes that can become infected.

The figure also includes a horizontal line illustrating a hypothetical the level of infections at which the healthcare system becomes saturated. Using the SIR model can induce us to overstate the negative effect of infections by over-predicting the speed at which the disease spreads and surpasses the healthcare system’s capacity to deal with the disease. Because the NET model predicts a flatter curve, the disease can actually be manageable when the stability of connections is taken into consideration. This is an important lesson.

Vaccination strategies: Because the dynamics of the two models are different, the efficacy of vaccination strategies will not be the same. Similar to Boppart et al. (2021), I consider a situation where a planner vaccinates a given percentage of the population, at random, as soon as the pandemic starts. Vaccinated agents become immune for the remaining periods.

Figure 8 shows the evolution of infected, susceptible, and recovered individuals in the network model (NET) without vaccines (left panel) and with vaccines (right panel). The drop in the number of susceptible individuals arises because I assume that 20% of them receive the vaccine, at once, on day 10. The policy is effective in slowing down the epidemics (flattening of the dashed-orange infections’ curve), while also reducing the extent of the epidemics: under the vaccination policy a significant fraction of the population never gets infected.

In order to compare the effects of vaccines between SIR and NET, I display the




Figure 7: Infection dynamics in NET in the benchmark (left) and with 20% vaccination (right)

evolution of infections in the SIR model with (dashed green line) and without (solid red line) 20% vaccination in the left panel of Figure 8. The right panel shows the comparison for the NET model. Qualitatively, vaccines have similar effects in both models. Quantitatively, they are more effective in a network model, particularly in slowing down the pandemic. Figure 9 shows infection dynamics in the network model for alternative vaccination rates. In this simple example, when 60% or more of the population is vaccinated, the infection never takes off. This would not happen in a SIR model, where local herd immunity is less likely. The second important lesson is that vaccination strategies tend to be more effective in network models.




Figure 8: Effects of vaccination on SIR (left) and NET (right) models




Figure 9: Infection curves under alternative vaccination levels

5 Conclusion

The comparisons above illustrate some important differences between infection dynamics in standard SIR and network models. We learned that because network models are characterized by more stable relationships between individuals, infections tend to travel more slowly (e.g. the curve is flatter) and vaccination policies tend to be more effective. In Boppart et al. (2021), the ‘fear factor’ is the force that endogenously flattens the curve. I have completely abstracted from it in this example, by assuming that individuals do not change the nature of their relationships as infections rise. Because of this, one cannot draw quantitative policy implications that are directly comparable to the original model studied by the authors. However, it points to a potential fruitful avenue of research by combining some elements of both models. Network models tend to abstract from the endogenous intensity of contacts, whereas economic models with SIR components—such as Boppart et al. (2021),—tend to abstract from network dynamics. ‘Second generation models’ could embed Boppart et al. (2021) into a network model. In addition to the dimensions discussed above, this would also allow us to consider targetted vaccination strategies that depend not only on age, but also on occupation. For example, it would make sense to vaccinate individuals working in the B sector (who tend to have less stable contacts form a wider base) at a faster rate than V -sector workers (who tend to meet the same team of co-workers every day).

References

Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi, “Spread of (mis)information in social networks,” *Games and Economic Behavior*, 2010, 70 (2), 194 – 227.

- Azzimonti, Marina, Alessandra Fogli, Fabrizio Perri, and Mark Ponder**, “Pandemic Control in ECON-EPI Networks,” Working Paper 27741, National Bureau of Economic Research 2021.
- **and Marcos Fernandes**, “Social Media Networks, Fake News, and Polarization,” Working Paper 24462, National Bureau of Economic Research March 2018.
- Boppart, Timo, Karl Harmenberg, Per Krusell, and Conny Olsso**n, “Integrated epi-econ assessment of vaccination,” Working Paper, IIES 2021.
- Fogli, Alessandra and Laura Veldkamp**, “Germs, Social networks and Growth,” *Review of Economic Studies*, 2020, *Forthcoming*.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and Jose-Victor Rios-Rull**, “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic,” Working Paper 27046, National Bureau of Economic Research 2021.
- Hur, Sewon**, “The Distributional Effects of COVID-19 and Optimal Mitigation Policies,” Working Paper 400, Federal Reserve Bank of Dallas 2021.
- Watts, D.J and S.H Strogatz**, “Collective dynamics of ‘small-world’ networks,” *Nature*, 1998, 393.