

24^e Congrès des économistes

What Can We Learn From Economico-Epidemiological Models?

Jean-Philippe Platteau (UNamur) & Marius Gilbert (ULB)

Epidemiologists use certain models (SIR or SEIR models, for example) to describe the natural evolution of an epidemic or a pandemic. The parameters of these models are based on various assumptions regarding biological and social phenomena. In contrast to this approach, economists consider human actions as the outcome of optimizing behavior based on private costs and benefits. The corresponding literature can be divided into two types depending on whether individuals are assumed to operate freely in a decentralized manner, or are subject to the public prescriptions of a central decision-maker acting as a benevolent social planner. This paper explores this literature and provides various examples showing, in particular, that this economic approach can help to explain part of the difference between the number of confirmed COVID-19 infections and the number predicted by epidemiological models.

General Considerations

Two features of the economic approach to human behavior need to be highlighted:

1. Economists considering human behavior not as a given but as an optimal response to a set of given circumstances. Such responses are typically driven by a rational calculus in the sense that individuals are assumed to choose the course of action that provides them maximum satisfaction given the constraints (budget or time constraints, for example) that they face. They are thus equipped with preferences (assumed to obey a number of properties ensuring that stable preference ordering is possible) and confronted with situations of scarcity, scarcity being understood as implying constrained optima.
2. Whereas in (neo)classical models of human behavior the behavior of other individuals was considered to be given, this is no more the case in modern, game-theoretic models. In the latter, individuals act strategically, as « chess players » - they hold beliefs regarding what others are going to do in a particular situation and they then act upon their beliefs. An equilibrium is achieved when beliefs and actions are mutually compatible and therefore reinforce each other: I can see that other people have behaved the way I thought and, consequently, my choice of action was right so that I have no reason to deviate from it. When everybody can reason thus, a so-called Nash equilibrium is reached.

There are important direct implications of the above features. Two of them deserve to be especially emphasized. First, the notion of constrained choice means that behind human decisions lie difficult trade-offs. For example, in a situation of epidemy, there is a trade-off between a health and an income objective: by isolating from potentially infected others, an individual reduces his/her risk of infection yet, simultaneously, s/he forsakes incomes that were otherwise accessible. Second, strategic thinking may easily yield multiple equilibria. For example, depending upon the prevailing beliefs, everybody can decide to wear a mask, but it is also possible that nobody does so. Equally possible are situations where a unique equilibrium prevails, that is, one where either conduct is established. Thus, if the situation requires that everybody wears a mask, it is possible that, left free to decide, individuals decide not to

wear a mask. We then say that the equilibrium which results from the decentralized free choices of the individuals is not socially optimal: individuals freely choose a bad outcome.

It bears emphasis that the notion of externality, one of the pillars of the theory of market imperfections, has been elaborated outside the framework of game theory. In the case of a negative externality, the individual does not take into account the adverse effects of her/his action on the others, simply because s/he does not have to bear the costs associated with these effects. At work here is classical individual rationality: in the absence of altruism, the individual acts in the light of her/his own interests, without consideration for the way others are affected by her/his own decisions. With strategic rationality, by contrast, the individual *may* take a decision that goes against her/his own interests, based on her/his beliefs and expectations regarding others' behavior.

Applications to epidemiological problems

The basic model used by epidemiologists, the so-called SIR model, distinguishes different states between an individual may move. The population is divided into three groups (compartments) of individuals: S , I and R (with $S+I+R=N$, where N is the population size). Group S is the group of susceptible individuals (i.e. those individuals that are at risk of being contaminated). At the beginning of the epidemic, S is the entire population given that nobody has anti-bodies (it is indeed a new virus for which no vaccine is available). Group I is the group of individuals who have been contaminated recently and are infectious. Finally, R is the group of individuals who were contaminated but had an outcome in the form of either a recovery or a death. They are not infectious anymore. The sizes of these groups evolve over time as the virus spreads. The size of S decreases when people become contaminated and move into the infectious group I . When individuals recover or die, they transition from the infectious group I to the group R . A system of 3 differential equations models the evolution of the sizes of these groups:

$$dS/dt = -\beta \cdot S \cdot (I/N)$$

$$dI/dt = \beta \cdot S \cdot (I/N) - \gamma \cdot I$$

$$dR/dt = \gamma \cdot I$$

The first equation states that the size of S , the number of susceptibles, decreases by the number of newly contaminated individuals, which is simply the transmission rate (β) multiplied by the number of susceptible individuals who came into contact with infectious individuals (I). More precisely, each susceptible person contacts β people per day, a fraction I/N of which are infectious.

The second equation states that the number of infectious individuals (I) will be increased by the newly contaminated individuals minus the previously infectious individuals who had an outcome and moved to group R (i.e., the removal rate, γ , multiplied by the infectious individuals (I).

The third equation states that the removed group increases by the number of individuals who were infectious and had an outcome (γI).

However, since there may be a significant incubation period for infections, as typically observed for Covid-19, there exists a latency period during which individuals have been infected but are not yet

infectious. To take this possibility into account, the SIR model is generally augmented of group E (for the exposed) and a new differential equation is added to the three equations of the basic SIR model to yield what is known as the SEIR model. The set of equations that define the SEIR model is:

$$dS/dt = -\beta \cdot S(I/N)$$

$$dE/dt = \beta \cdot S(I/N) - \sigma \cdot E$$

$$dI/dt = \sigma \cdot E - \gamma I$$

$$dR/dt = \gamma I$$

where σ is the incubation rate, i.e., the rate at which latent individuals becoming infectious. The increase in the number of infectious people (see the third equation) is now the number of contaminated people who have reached the end of their latency period minus the previously infectious individuals who moved to group R .

By simulating this type of models, epidemiologists are able not only to describe the evolution of an epidemic as it unfolds naturally, but also to predict the impact of various policy measures adopted by public authorities to combat it. These models therefore offer an excellent starting point for both analysis and policy thinking. It is evident, however, that such simulation exercises necessarily involve a lot of assumptions about key parameters that are both biological and social. As examples of the former, we have the duration of the latency period; the probability of exhibiting symptoms; the probability of being contaminated, when in contact with either an asymptomatic or a symptomatic individual; the infectious period for symptomatic and asymptomatic individuals. As examples of the latter, we can think of frequency of contacts between individuals, preferably differentiated by age classes, or individual reactions to various public health measures. In regard of social parameters, transitions in epidemiological models are determined by aggregates whose behavior is decided in a rather ad hoc manner. For example, the values of cells in a contact matrix are assumed to be constant at least for a given period of time. If they are allowed to vary, say to take account of a Covid fatigue that leads individuals to relax their contact discipline, the adjustments of these values are made in a discontinuous manner and based on rules of thumb or, in the best cases, on panel data collected for the purpose. Likewise, how individuals respond to interventions such as travel restrictions, school closures, bans on meetings (private and/or public), and quarantine prescriptions (after returning from a red zone abroad or after being tested positive) is modeled by using guesses about compliance rates.

More recent versions of the SIR model, however, make contact rates depend on the heterogeneous topology of the network of contacts and mobility of people across locations, or they make the infection rate depend on the activity intensity of each node of the network (see Alfaro et al., 2020 for references). Thus, in the so-called SIR-network models, different groups are distinguished that have different exposure or contact rates to each other. These groups are usually defined on the basis on the intensity of their internal contact rates, the underlying idea being that these rates tend to be higher among peer groups (e.g., age groups, or people meeting on the workplace).

In contrast to the above approach, economists want to give pride of place to the role of incentives, which requires a foundation in individual decision-making. As they see it, human actions are the outcome of optimizing behavior based on an evaluation of private costs and benefits. And when successive periods of time are considered, such as must be the case in the analysis of decisions under an epidemic,

optimization rests on an assessment of the costs and returns of doing something now against the expected future payoffs (Garibaldi et al., 2020, p. 2).

As expected, there has been a recent surge in economics papers dealing with the Covid crisis, including papers with a theoretical approach to several challenging issues raised by the epidemic. At the heart of a significant fraction of these theoretical papers are decision problems that individuals must solve by themselves. This new literature can be divided into two types depending on whether individuals are assumed to operate freely in a decentralized manner, or are subject to the public prescriptions of a central decision-maker acting as a benevolent social planner. In the first case, the role of the government is limited to announcing the epidemic and a key question is how the decentralized equilibrium thus attained can be compared to the equilibrium achieved in a standard epidemiological model as well as to the social optimum. In the second case, the role of the government goes much further and consists of devising and implementing a set of policy measures. Here, the focus is on the characteristics of an optimal lockdown policy and its evolution through time. This important distinction is adopted in the short survey that we now present.

A short survey of theoretical contributions by economists

Individual behavior following the announcement of an epidemic

There are various decisions that individuals have to make when facing an epidemic. In particular, they must decide how much they want to interact with others, and whether they will preventively protect themselves, by using preventive protection equipment (masks, face shields, gloves), adopting hygienic measures (hand sanitizers), and/or practicing physical distancing in human contacts. These two decisions can be analyzed jointly or separately.

Analyzing them jointly is done in Bhattacharya et al. (2021). They consider agents who have to choose whether to socialize or self-quarantine, and whether to use preventive protection. Self-quarantine implies a fixed loss of income units and this cost is disproportionately high for the poor. People can be in three different health states represented by different proportions in the population: asymptomatic-healthy (AH), asymptomatic-infected (AI), and symptomatic-infected (SI). Moreover, since protection is costly and does not bring any benefit to an agent under self-quarantine, only the protection decision for agents who have chosen to socialize needs to be examined. In making this decision, an agent has to take into account s/he may be either AH or AI type and s/he can be infected only if of the former type. The chance of getting infected depends on meeting an AI agent whose protection choices s/he does not know ex ante. Conditions for the following combined decisions are derived: self-quarantine, unprotected socialization, and protected socialization.

In such a setup, expectations play a key role in determining outcomes. For instance, if an agent perceives the risk of infection to be high because of an anticipation that the proportion of socializers who are unprotected is high, s/he will have higher incentive to use protection. Moreover, the two decisions of socializing and protecting are interlinked: while reducing costs, less use of preventive protection raises the risk of infection during social interactions, but it may simultaneously incite more people to stay at home because of a higher fear of being infected by others.

Equilibrium analysis yields the conclusion that different equilibria can arise depending on the values of the model's parameters: if a positive fraction of agents of either the AH or the AI type always choose to

self-quarantine, it is possible that all socializing agents go protected, all go unprotected, or a fraction of them go protected. A merit of the exercise conducted in this paper is that the authors allow for agent heterogeneity. This allows them to show that the trade-off between cost and benefits of the two prevention choices, social isolation, and the use of preventive protection equipment varies across the income distribution. More precisely, a (mean-preserving) increase in pre-existing income inequality unambiguously increases the equilibrium proportion of unprotected, socializing agents and may increase or decrease the proportion who self-quarantine.

Many papers focus on social distancing decisions alone. In this strand of the literature, the contact rate of the epidemiological model is considered to result from a decision problem on the extent of social interactions. The formal approach is based on the idea that an individual maximizes an expected utility that depends with probability p on the possibility of being contaminated and the probability $1-p$ of being in good health. The probability p is itself a function not only of the amount of social activity chosen by the individual himself or herself but also on the amount chosen by the other susceptible or infected individuals (the average amount of social activity). It is this latter assumption that sets the strategic framework adopted for the analysis: an individual takes a decision that other individuals take at the same time and the equilibrium is defined in such a way that all these decisions are compatible with each other.

A major result here is the following: when people freely choose the amount of social distancing, that amount is lower than what is socially optimal (that is, the level a social planner would select). Inefficiency of the decentralized equilibrium is clearly due to an externality: when choosing their own social activity, people ignore the infectious impact that their social interactions have on others and they therefore decide to go out too much. In other words, people think of the risk of being infected by other people but not of the risk of contaminating them (Caulkins et al., 2020; Chang and Velasco, 2020; Farboodi et al., 2020; Alfaro et al., 2020; Eichenbaum et al., 2020). Moreover, they ignore the congestion externalities that expose the available medical facilities to the risk of acute stress (Boucekkine et al., 2020; Garibaldi et al., 2020; Ichino et al., 2020). Finally, there exists a dynamic externality (called immunity externality) the effect of which goes in the opposite direction: restricting social contacts retards herd immunity (Garibaldi et al., 2020).

Chen (2012) has refined the analysis through the study of the role of the contact function that governs the rate at which encounters occur in public. He shows that the result of the comparison between decentralized equilibria and the social optimum critically depends on this contact function. If the contact ratio is increasing in the number of people out in public, then there exists a unique (Nash) equilibrium which differs from the social optimum: in this case, the amount of public avoidance is too low from a social welfare point of view. If the contact ratio is decreasing, there can exist multiple equilibria, none of which is in general socially optimal. Finally, if the contact ratio does not vary with the number of people out in public, there is a unique Nash equilibrium and it is also the socially optimal outcome.¹

In comparing outcomes obtained under the traditional SIR model of the epidemiologists with a decentralized epidemic equilibrium driven by rational forward-looking agents, Garibaldi et al. (2020) have paid attention to herd immunity. They find that when agents optimally decide their level of social interactions, the longer time they need to reach herd immunity comes with a large gain in the form of

¹ Society's welfare is the sum of the utility of all agents in the population, i.e., the utility of all infected agents + the utility of all recovered agents + the utility of all susceptible agents. The choice of public activity by susceptible agents does not affect the utility of infected and recovered agents. As a result, the problem of maximizing social welfare in a given period is equivalent to one of choosing the susceptible agents' public activity level that maximizes the utility of all the susceptible agents, i.e., the socially optimal outcome.

avoiding illness among a substantial fraction of the population. Moreover, the number of people who get infected before herd immunity is reached is much lower than the number obtained in the standard SIR model. Along the decentralized epidemic equilibrium, optimal social activity clearly follows a U-shaped behavior: the level falls as the epidemic spreads, then reaches a minimum before it starts rising until the steady state (that is, an equilibrium trajectory in which all variables evolve at a constant rate).

Toxvaerd (2020) confirms that in a decentralized equilibrium, the aggregate level of infection across the epidemic is lower than what a traditional non-economic epidemiological analysis would suggest, thus indicating that the latter overstates the severity of the epidemic by ignoring rational human responses. Furthermore, uncoordinated social distancing acts to flatten the curve of the epidemic by reducing peak prevalence. And if, in equilibrium, it stops once herd immunity sets in, it nevertheless acts to extend the duration of the epidemic beyond the benchmark of a non-behavioral epidemiological model. It is a striking result that the comparative-static predictions of the economic model are the reverse of those in the uncontrolled epidemiological model.² For example, peak prevalence and cumulative incidence are both increasing in the infectiousness of the disease in the epidemiological model, whereas they are decreasing in the economic model. The rationale is that the endogenously determined social distancing decisions of the individuals react to higher infectiousness by engaging in more protective behavior. Finally, the epidemic curve becomes flatter in the economic model not only as the disease becomes more infectious, but also as the health consequences of the disease become more severe for the individuals.

Figure 1 (Our World in Data, 2020) shows the significant discrepancy between the estimated Covid-19 infections and the officially confirmed cases in Germany (with public intervention). The estimations are based on the IHME (Institute for Health Metrics and Evaluation) model, “a hybrid with two main components: a statistical ‘death model’ component produces death estimates that are used to fit an SEIR model component” (Covid-19 Projections, 2020).³

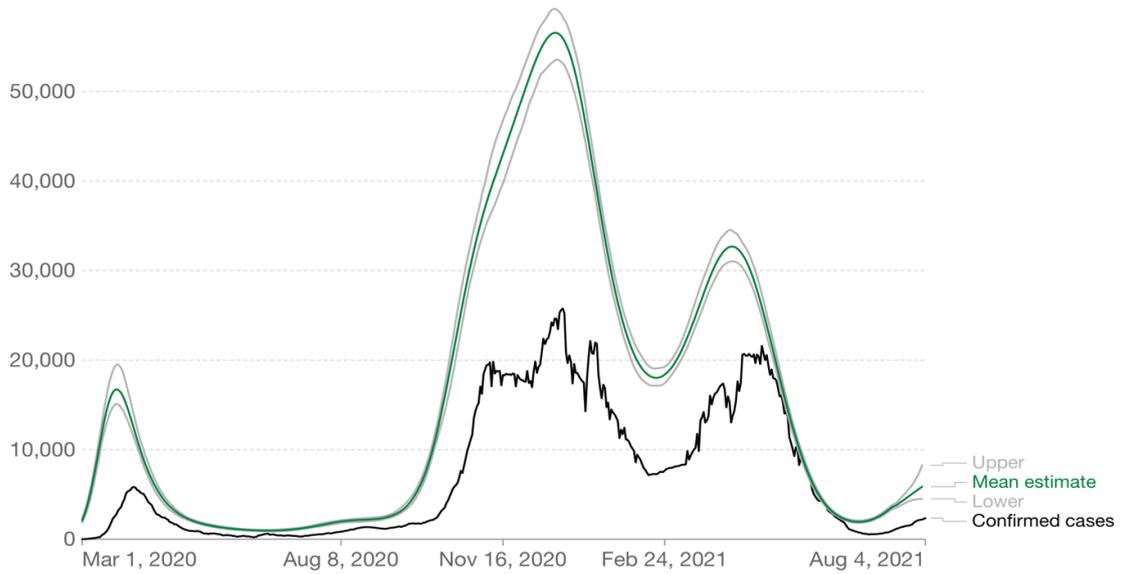
² By a comparative-static result, economists mean the effect, all else being equal, of the variation of a model’s parameter on an endogenous variable at equilibrium.

³ The issue here is how are confirmed cases observed. To have a valid comparison, we would need them to be accurately observed. Unfortunately, this is never the case as we would have needed a large random testing of the population, which was never organized. Since there are many asymptomatic cases that escape detection through doctors and hospitals, a huge bias exists. Could you therefor say more about the way confirmed cases are computed?

Figure 1: Estimated daily Covid-19 infections in Germany

Daily new estimated COVID-19 infections from the IHME model, Germany

Estimates of the true number of infections. The "upper" and "lower" lines show the bounds of a 95% uncertainty interval. For comparison, confirmed cases are infections that have been confirmed with a test.



Source: IHME (2021), JHU (2020)

Note: This chart shows the model estimates dated 4 August 2021.

OurWorldInData.org/covid-models • CC BY

Source: Our World in Data, 2021

Observe that the IHME model produces overestimates whenever daily infections are increasing or assume their peak values. Obviously, the model does not adequately take into account precautionary measures of the individuals. Epidemiologists sometimes describe this observation as “prediction paradox”.

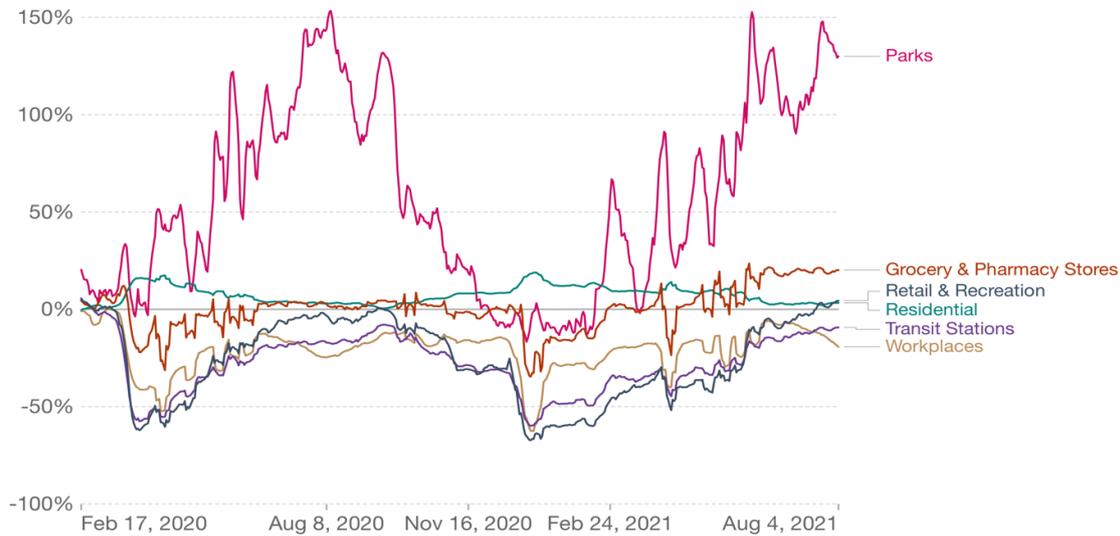
The situation is somewhat different whenever daily infections decline. Individuals tend to reduce their precautionary measures and the estimated values get closer to the actual values.

Figure 2: Changes in the number of visitors in Germany in the course of the pandemic

How did the number of visitors change since the beginning of the pandemic?, Germany



This data shows how community movement in specific locations has changed relative to the period before the pandemic.



Source: Google COVID-19 Community Mobility Trends – Last updated 21 August 2021, 15:53 (London time)
Note: It's not recommended to compare levels across countries; local differences in categories could be misleading.
OurWorldInData.org/coronavirus • CC BY

Source: Our World in Data, 2021

An interesting feature of Toxvaerd's model is that it represents individual social distancing decisions as the engine of a flow rate regulation between healthy and recovered individuals, where the underlying uncontrolled flow rates are determined by the biological features of the disease. The mechanism works as follows. For sufficiently low disease prevalence, say at the beginning or the end of an epidemic, the risks from social interactions are small and individuals therefore choose not to socially distance themselves. For higher levels of disease prevalence, the risk of exposure may outweigh the benefits so that individuals are incited to switch to social distancing. Aggregate disease prevalence at equilibrium remains constant through time until a sufficiently high number of individuals have transited through the cycle S-I-R to cause disease prevalence to fall without further social distancing. As in Garibaldi et al. (2020), social activity follows a U-shape pattern (or, social distancing follows an inverted-U shaped pattern).

Figure 2 (Our World in Data, 2021) shows how the number of visitors changed in Germany in the course of the pandemic. Observe that – as predicted in the economic models – social activities were high when daily infections were low, and vice versa. Again, these observations include public policy measures.

Baril-Tremblay et al. (2021) have adopted a rather similar perspective to address the problem of how much time individuals strategically choose to spend interacting with others. They show that, when self-isolation is costlier than being sick, agents do not self-isolate in equilibrium and the dynamics of the epidemic is the same as in the SIR model. When the opposite is true, however, and the cost of confinement is relatively small, the (symmetric) equilibrium may be such that agents partially self-isolate at each date. In addition, this equilibrium is non-stationary. Population reacts to the epidemic announcement by self-isolating drastically, which results in a drop in the proportion of people infected. Then, agents gradually increase the time they spend outside, and the effective reproduction number is maintained below unity. This last outcome is obtained because agents compensate the decrease in the

risk of infection by reducing social distances, yet not to the point of accelerating the epidemic. Unlike the bell-shaped curve of the SIR model, the epidemic curve decreases between the time of announcement and the arrival of the vaccine.

Social distancing may be achieved not only through the restriction of social encounters but also through the use of preventive protection devices, such as masks. Ng (2021) has precisely studied the mask-wearing behavior of individuals confronted with an epidemic. He assumes that they know that a mask protects people around the wearer more than it protects the wearer herself (himself). The central question he addresses is whether mask-wearing behaviors are discouraged by free riding or mutually reinforced by strategic complementarity. The overall intuition is that whenever the cost of wearing a mask is sufficiently low, it is in the interest of everyone to wear a mask, and vice-versa: if this cost is sufficiently large, everyone will refuse to use the mask. How low the cost has to go below which universal mask-wearing would be voluntarily adopted by the population depends on the following other key factors: (i) the fraction of people expected to have been infected without showing symptoms; (ii) the filtration efficiencies of the masks; and (iii) the number of people whom an individual comes across in normal life (proxied by population density). Population density plays a pivotal role because the risk of infection increases with it, thereby raising the profitability of wearing a mask.

If population density exceeds a given threshold, then two (expectational) equilibria are possible: everyone wears a mask or nobody does.⁴ It is precisely in this instance that a mask mandate can make sense since it may tilt the system from the bad to the good equilibrium. Since increasing the others' probability of wearing masks then moderates the infection risk, incentivizing an individual to also wear a mask to stay healthy works for the benefit of everybody. Otherwise, a mask mandate is of no use. It would be either socially inefficient –the cost is so high that it will never be in the interest of the society to wear a mask–, unnecessary –there is a unique (Nash) equilibrium in which everyone wears a mask–, or incentive-incompatible –nobody wants to wear a mask at another unique (Nash) equilibrium. In the latter case, a mask mandate is doomed to yield low compliance. In other words, it is only when population density is high enough (say, in relatively crowded cities) and the cost of mask-wearing is low enough (say, in circumstances where people have no cultural resistance against wearing a mask) that a public intervention is desirable to defend the general interest. In such circumstances, indeed, although it is socially beneficial for everyone to wear a mask, decentralized individuals may lack the private incentives to do so (p. 68).

Note finally that economics is often portrayed as a “dismal science” because it is grounded on the assumption of selfish individuals. Yet, this assumption can be relaxed and the consequences of this relaxation can be investigated. Thus, Alfaro et al. (2020) assume that infected individuals hold some altruistic preferences, implying that since they partly internalize the risk of infecting susceptible individuals, they chose a reduced level of social activity (in their model, there is no distinction between symptomatic and asymptomatic individuals).

Public interventions: the optimal lockdown

An important lesson from the preceding section is that, because people fear infection, the impact of public interventions, such as a lockdown, will be smaller than generally predicted on the basis of models that ignore endogenous human responses to the simple announcement of an epidemic (see Figure 1, in

⁴ This happens when expectations play a key role. If an individual expects that all other individuals will wear a mask, her/his interest is in also wearing a mask because the cost of doing so is smaller than the private benefit. If, on the other hand, s/he expects the others to abstain from mask-wearing, her/his interest is again to follow suit because the cost of wearing a mask now exceeds the benefit.

particular). On the other hand, a lockdown is useful because a decentralized equilibrium outcome is generally suboptimal. The public nature of the decisions at stake, such as self-isolation and mask-wearing, is what invites free riders and causes individual rationality to possibly lead to collective irrationality. It is therefore not surprising that economists have devoted much effort to analyzing the characteristics of an optimal lockdown policy. This is the so-called social planner problem. Since this literature is rather abundant, lack of space constrains us to limit ourselves to extracting just a few interesting lessons from a selected number of papers.

Garibaldi et al. (2020) show that a benevolent social planner who maximizes the overall welfare of the population while taking account of various types of externalities, will choose a lower level of social activity than what the individuals themselves would privately choose.⁵ Compared to the decentralized equilibrium, the planner thus helps reduce infections and the pressure on medical facilities. In a dynamic setting, however, things are more complicated: if the medical externalities are expected to be more important in the more distant future, the planner on the margin may prefer that a higher number of people fall ill early on (when there is spare capacity in the health sector) rather than later (when the capacity constraint binds). Also, when comparing the private and social equilibria, only the herd immunity externality provides incentives to the planner to speed up the spread of the epidemic. Yet, the possibility of obtaining a vaccine in the future reduces the positive externality associated with a higher number of recovered individuals (pp. 12-13).

A major contribution to the study of optimal lockdown has been provided by Caulkins et al. (2021) who use an optimal control model. They ask two questions: what is the optimal intensity with which to lockdown, and how should that intensity vary dynamically over the course of an epidemic? In addressing these questions, the constraint raised by limited medical facilities receives primary attention. The analysis concludes that several broad strategies emerge: they range from brief lockdowns that only “smooth the (infection) curve” to sustained lockdowns that prevent infections from spiking beyond the healthcare system’s capacity. It can even be optimal to have two separate periods of locking down so that returning to a lockdown after initial restrictions have been lifted should not necessarily be considered as a sign of failure. In addition, the authors find that relatively small changes in judgments about how to balance health and economic harms can alter dramatically which strategy prevails. This is because there are parameter configurations for which two or even three of the distinct optimal strategies can all perform equally well for the same set of initial conditions. The implication is that even people who share a common understanding of the problem’s economics and epidemiology can prefer dramatically different policies.

The public policy responses can be described by the “Government Stringency Index”, which is a composite measure of the strictness of policy responses. It includes school and workplace closures, restrictions on public gatherings, transport restrictions, and also stay-at-home requirements (for details on the “Oxford COVID-19 Government Response Tracker” see Ritchie et al., 2020).⁶

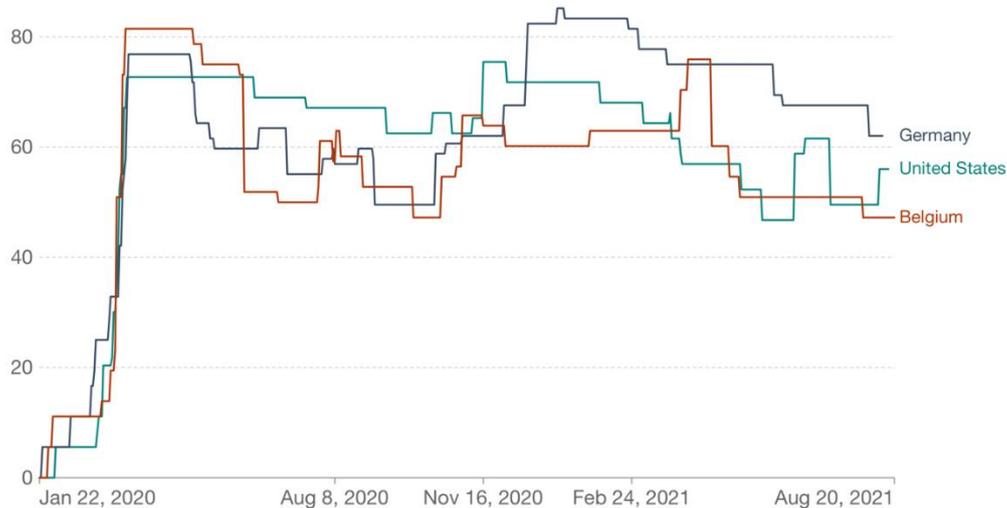
⁵ Technically, the authors assume that the social planner is aware that the equilibrium is a symmetric Nash equilibrium and that contacts involve at least two individuals (p. 9).

⁶ Note that Belgium had quite strict regulations regarding private meetings: the number of people you could invite at your home was very limited, and they had to be always the same persons. This is not taken into account in the Oxford Index. Of course, there is the problem of the enforceability of such measures, which is quite low. Still, they served as a signal of the seriousness of the pandemic.

Figure 3: Government Stringency Index for Belgium, Germany and the United States

COVID-19: Stringency Index

This is a composite measure based on nine response indicators including school closures, workplace closures, and travel bans, rescaled to a value from 0 to 100 (100 = strictest). If policies vary at the subnational level, the index is shown as the response level of the strictest sub-region.



Source: Hale, Angrist, Goldszmidt, Kira, Petherick, Phillips, Webster, Cameron-Blake, Hallas, Majumdar, and Tatlow (2021). "A global panel database of pandemic policies (Oxford COVID-19 Government ResponseTracker)." *Nature Human Behaviour*. – Last updated 21 August 2021, 20:50 (London time)
OurWorldInData.org/coronavirus • CC BY

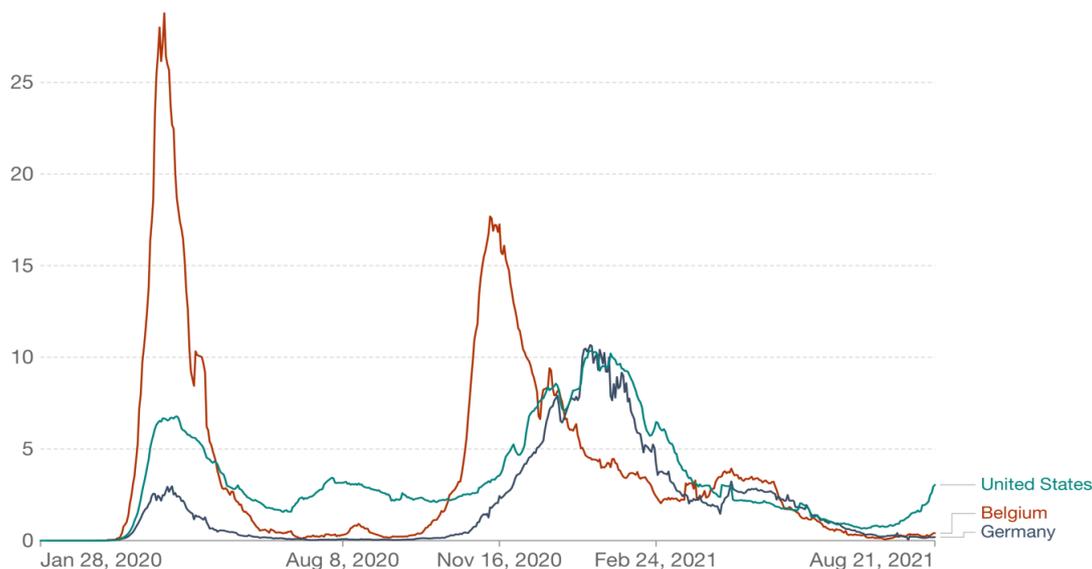
Source: Our World in Data, 2021

Figure 3 (Our World in Data, 2021) details the public policy responses for Belgium, Germany and the United States in the course of the pandemic. According to this figure, there are only small differences between the policies responses of the three countries. More stringent regulations seem to correspond somewhat to the new confirmed deaths, as indicated in Figure 4. In this sense, we obtain an interesting chain of relationships: Covid-19 cases and deaths affect policy responses and individual protective behaviors, which affect the numbers of cases and deaths. Nevertheless, the changes in the policies are limited, given the drastic variations in the number of cases, for example. Figure 5 (Our World in Data, 2021) shows the number of daily cases.

Figure 4: Daily new confirmed COVID-19 deaths in Belgium, Germany and the United States

Daily new confirmed COVID-19 deaths per million people

Shown is the rolling 7-day average. Limited testing and challenges in the attribution of the cause of death means that the number of confirmed deaths may not be an accurate count of the true number of deaths from COVID-19.



Source: Johns Hopkins University CSSE COVID-19 Data

CC BY

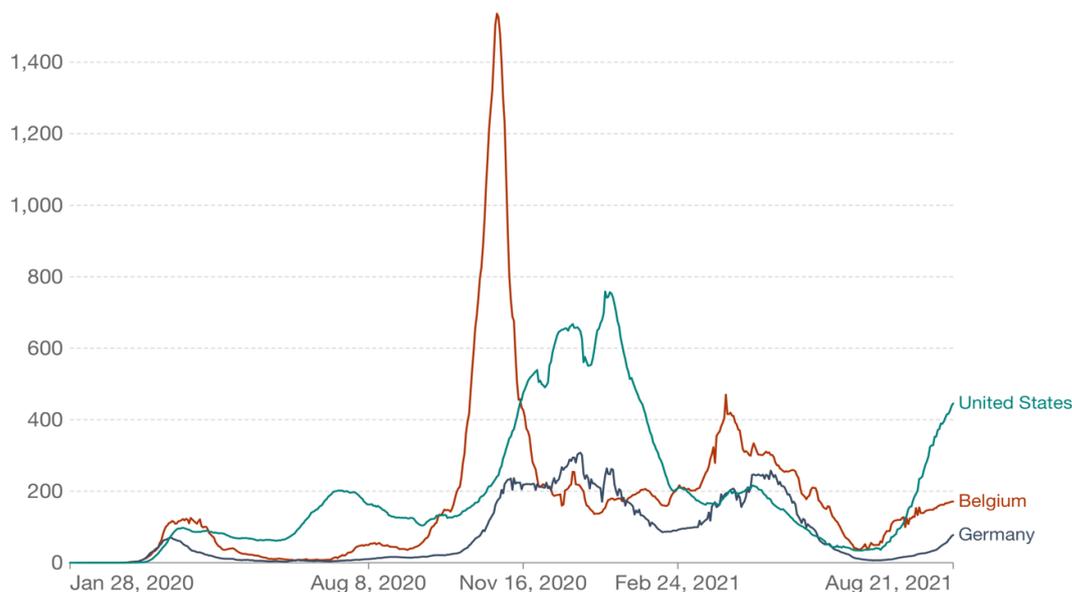
Source: Our World in Data, 2021

Another paper whose authors center their attention to the role of medical externality is Loertscher and Muir (2021). It attempts to determine the minimal lockdown that satisfies the constraint set by the maximum healthcare system capacity. The epidemiological model is augmented by an economic production function to analyze tradeoffs involving economics. These tradeoffs result from the fact that the amount of productive labor is assumed to be inversely related to the severity of the lockdown chosen by the policy-maker. With homogeneous agents, the conclusion is easily reached that, all else equal, states or countries with larger healthcare capacities can afford less stringent lockdowns. The situation gets more complicated when the population is composed of heterogeneous agents corresponding, say, to different age classes. Here, the authors assume that there is a continuum of types, the type of any given individual is observable, and the policy-maker can implement a type-dependent lockdown policy that consists of either a zero or a complete lockdown. They now find that the optimal policy is bang-bang: there is a critical threshold type so that zero lockdown is optimal below it (say, for classes younger than a certain age) and complete lockdown above it (for the older classes). This is true whether mixing is homogeneous among all type cohorts or type-dependent.

Figure 5: Daily new confirmed COVID-19 cases for Belgium, Germany and the United States

Daily new confirmed COVID-19 cases per million people

Shown is the rolling 7-day average. The number of confirmed cases is lower than the number of actual cases; the main reason for that is limited testing.



Source: Johns Hopkins University CSSE COVID-19 Data

CC BY

Source: Our World in Data, 2021

As the authors point out, since the policy-maker maximizes economic output subject to a medical capacity constraint, the optimization problem has the advantage of avoiding the difficult decision about how economic activity is traded off against the number of deaths caused by the epidemic. There is thus no need to specify the value of human life (p. 8). Another interesting feature of the proposed model is that once the capacity constraint becomes slack, no future policy interventions are required. This is because the optimal dynamic policy leads to the shortest possible duration of the lockdown by decreasing the population of susceptible individuals as efficiently as possible, subject to the constraint. In principle, a second wave of infections cannot occur. A statically optimal policy would imply a longer and less severe lockdown, resulting in an extended period of depressed output (p. 9). Moreover, “even after the peak of the pandemic has passed, if a policy-maker cancels a statically optimal lockdown too soon, this can result in a large second wave of infections occurring”. This is more likely to happen if the capacity constraint is tight and a large population of susceptible individuals remain after the peak. In this sense, write the authors, “statically optimal policies are less robust to future mistakes on the part of policy-makers” (pp. 9-10).

Unlike in the previous exercises, Fernichel (2013) examines how the efficiency of a lockdown can be impaired by a lack of information. The point of departure of his analysis is the idea that otherwise identical individuals in different health classes face different incentives and therefore behave differently. He explores two scenarios. In the first scenario, the social planner is in full control in the sense of being able to provide targeted incentives across health classes, such as assumed in Acemoglu et al. (2020). He chooses the contact levels for susceptible, infected and recovered individuals, which means that he chooses behaviors directly. In the second scenario, the planner is constrained: public social distancing policies are not sufficiently flexible to provide targeted incentives across health classes. As a consequence, policies destined to encourage social distancing are blunt and provide incentives for all individuals to reduce contacts. A salient conclusion is that the decisions of a constrained social planner

can make the society worse-off than decentralized decision-making. An implication is that the oft-neglected behavior of recovered and immune individuals is important for welfare and health outcomes.

A different type of complexification of the social planner's problem has been introduced by Bandyopadhiyay et al. (2021). It consists of allowing for habit formation in individual behavior. The central result obtained by the authors is that an early lockdown can be beneficial not only to slow down the spread of an epidemic, but also to create beneficial formation of habits, such as social distancing and hygienic precautions.

Finally, Bosi et al. (2021) study the optimal lockdown policy in a dynamic general equilibrium model where households are altruistic in the sense that they feel empathy towards the infected individuals. They argue that without empathy the optimal lockdown policy chosen by the policy-maker is a zero lockdown (a government facing selfish individuals does not confine the population), while under empathy a positive lockdown is optimal. Moreover, the optimal lockdown is positive only beyond a critical degree of altruism, and its severity then increases in the degree of altruism prevailing in the population. The reason why zero lockdown is optimal in the absence of empathy (altruism) is simply that selfish agents do not value the state of health of other people. On the contrary, when agents are altruistic, a substitution mechanism enters into play: "households are willing to accept a lower consumption in exchange for healthier people" (p. 6). In both cases, however, because the economic and social costs of a lockdown are taken into account by the social planner, it is efficient to reach an endemic steady state with a positive share of infected people in the population. This is a major contrast with the conclusions derived from pure epidemiological models, which recommend to eradicate the epidemic as quickly as possible. In other words, "even the simplest model encompassing epidemics and economics finds there is a conflict between health and production, which only empathy can partially overcome" (p. 6).

Conclusion

The integration of behavioral equations in epidemiological models seems to bear valuable fruits. We learn that the trajectory of an epidemic looks less worrying when due account is taken of the self-disciplining human responses based on the fear of infection. At the same time, these spontaneous self-limiting behaviors are not sufficient compared to the restrictions which a central authority would impose in the name of collective welfare. Several types of externalities explain the divergence between the equilibrium outcome of a decentralized mechanism and the social optimum prescribed by such an authority. If the weaknesses of standard epidemiological models arise from their rather mechanical character, –human beings are absent since they do not genuinely act–, the economic-epidemiological models suffer from two severe limitations. First, agent heterogeneity is not sufficiently taken into account. For instance, some models allow for several age or health classes but fail to distinguish between people depending on whether they have a robust or a fragile health, or whether they are health-anxious or not (which are not exactly the same things). The main problem with introducing such complexities into economic-epidemiological models is that they would become sorts of black boxes preventing clear analytical lessons, which economists tend to avoid.

Second, social optimum models, by nature normative, do not say anything about the enforceability of the optimal policies. In the ideal case, people are civic: they obey to governmental prescriptions out of respect for the authority or simply because it is their duty to do so. Bearing in mind Bosi et al. (2021), this means that altruistic individuals are not necessary to effectively combat an epidemic. Selfish

individuals may do the job provided that they are civic, that is, they act in the light of their own private interests if left free to behave as they wish, yet they are ready to modify their behavior if asked by their government. Unfortunately, the hypothesis of civicness is not applicable to many countries, in particular those where the legitimacy of the government is weak for reasons that are historical (bad precedents in matters of public health management), political (low trust in authoritarian regimes that have proved incompetent or unreliable), or social (strong polarisation of the society). Then arises the question of the degree of coercion that the government wants to use and of the cost involved. If its legitimacy is so low that the cost of coercion would exceed the expected benefits, the constrained optimum is zero confinement. Therefore, countries whose political culture is comparatively civic thus have an obvious advantage in their struggle against an epidemic.

The problem gets even more complicated once we drop the idea of stable individual preferences. Thus, weariness can affect people subject to long periods of lockdown or to repeated lockdown episodes. Such weariness is unavoidably reflected in an erosion of civic norms with the effect of undermining the efficacy of public health policies. In this case, it may be preferable to implement a less severe lockdown even though it is not first-best optimal (i.e., an unconstrained optimum unhampered by enforceability problems). Here, modelling turns out to be difficult not only because of the very complexity of the problem but also because of the numerous assumptions required at the level of human behaviors and their dynamics and at the level of the constraining instruments available to the government. Among the latter are the various ways of punishing rule violations (including the setting of the fines) and decisions regarding the detection of these violations and the amount of resources to be devoted to the task.

References

Acemoglu, D, I Chernozhukov, M D Werning and V Whinston (2020), A Multi-Risk SIR Model with Optimally Targeted Lockdown, NBER Working Paper 27102.

Alfaro, L, E Faia, N Lamersdorf and F Saidi (2020), Social Interactions in Pandemics: Fear, Altruism, and Reciprocity, CEPR Discussion Paper, DP14716.

Bandyopadhyay, S, K Chatterjee, K Das and J Roy (2021), “Learning Versus Habit Formation: Optimal Timing of Lockdown for Disease Containment”, *Journal of Mathematical Economics*, 93, Special Issue on Covid.

Baril-Tremblay, D, C Marlats and L Ménager (2021), “Self-Isolation”, *Journal of Mathematical Economics*, 93, Special Issue on Covid.

Bhattacharya, J, S Chakraborty and X Yu (2021), “A Rational-Choice Model of Covid-19 Transmission with Endogenous Quarantining and Two-Sided Prevention”, *Journal of Mathematical Economics*, 93, Special Issue on Covid.

Bosi, S, C Camacho and D Desmarchelier (2021), “Optimal Lockdown in Altruistic Economies”, *Journal of Mathematical Economics*, 93, Special Issue on Covid.

Boucekkine, R, A Carvajal, S Chakraborty and A Goenka (2021), “The Economics of Epidemics and Contagious Diseases: An Introduction”, *Journal of Mathematical Economics*, 93, Special Issue on Covid.

Caulkins, J P, D Grass, G Feichtinger, R F Hartl, P M Kort, A Prskawetz, A Seidl and S Wrzacek (2021), “The Optimal Lockdown Intensity for Covid-19”, *Journal of Mathematical Economics*, 93, Special Issue on Covid.

Chang, R, and A Velasco (2020), Economic Policy Incentive to Preserve Lives and Livelihoods, NBER Working Paper N° w27020.

Chen, F (2012), “A Mathematical Analysis of Public Avoidance Behavior During Epidemics Using Game Theory”, *Journal of Theoretical Biology*, 278, pp. 107-119.

Eichenbaum, M S, S Rebelo and M Trabandt (2020), The Macroeconomics of Epidemics, NBER Working Paper, WP26882.

Farboodi, M, G Jarosch and R Shimmer (2020), Internal and External Effects of Social Distancing in a Pandemic, NBER Working Paper N° w27059.

Fenichel, E P (2013), “Economic Considerations for Social Distancing and Behavioral Based Policies During an Epidemic”, *Journal of Health Economics*, 32(2), pp. 440-451.

Garibaldi, P, E R Moen and C A Pissarides (2020), “Modelling Contacts and Transitions in the SIR Epidemics Model”, *Covid Economics*, CEPR (London), Issue 5, 16 April 2020, pp. 1-19.

Ichino, A, C A Favero and A Rustichini (2020), Restarting the Economy While Saving Lives Under Covid-19, CEPR Discussion Paper, DP14664.

Loertscher, S, and E V Muir (2021), “Road to Recovery: Managing an Epidemic”, *Journal of Mathematical Economics*, 93, Special Issue on Covid.

Ng, T (2021), “To Mask or Not To Mask”, *Covid Economics*, CEPR (London), Issue 81, 16 June, pp. 50-72.

Toxvaerd, F (2020), “Equilibrium Social Distancing”, *Mimeo*, Faculty of Economics, University of Cambridge, UK.

Ritchie, H, E Mathieu, L Rodés-Guirao, C Appel, C Giattino, E Ortiz-Ospina, J Hasell, B Macdonald, D Beltekian and M Roser (2020), Coronavirus Pandemic (COVID-19), Published online at OurWorldInData.org. Retrieved from: <https://ourworldindata.org/coronavirus> [Online Resource]