Confronting epidemics: the need for epi-econ IAMs

Timo Boppart*, Karl Harmenberg†, John Hassler‡,

Per Krusell§ and Jonna Olsson¶

December 10, 2020

Prepared for the National Institute of Economic Research

Abstract

We discuss what tools would be useful in confronting epidemics, especially from the perspective of economics. Our main proposal is for policymakers to employ "epi-econ IAMs": explicit Integrated Assessment Models, where epidemiology is integrated with economics. These models are under rapid development, but arguably not yet quite ready for quantitative use.

*Institute for International Economic Studies, Unversity of St. Gallen
†Copenhagen Business School
‡Institute for International Economic Studies
§Institute for International Economic Studies
¶University of Edinburgh

Contents

1 Introduction

The current covid-19 pandemic has necessitated difficult decisions for politicians and government officials across the world. In particular, these decisions are, quite directly, "life-and-death decisions". This document outlines what we believe would be a very useful tool for policymakers in such circumstances. The idea is for the tool to allow a systematic evaluation of, and comparison between, a variety of policy options. No policy that we are aware of can make the fundamental problem—the epidemic—disappear magically, nor is there a policy that dominates other policies in all dimensions. In short, policy making will be hard and painful, and having a systematic tool for comparison then becomes all the more useful. The kind of tool that we are advocating is currently in the process of being built and tested by literally thousands of researchers around the world, including ourselves. However, as we will argue, we also believe that more work is needed before the models can be put to concrete use.

The structure we propose departs from two presumptions. One is that we take for granted that there is a fundamental trade-off between, on the one hand, the loss of lives and, on the other, the loss of welfare for those alive. This presumption needs one qualification: it appears conceivable in some cases, early on in an epidemic, to detect, trace, and isolate very quickly and effectively, so as to essentially eradicate the virus. The cases of South Korea or Taiwan have, in fact, been described in these terms. In other words, our focus here is on a situation where the window of opportunity for such a solution has closed (or was never available).[1] However, we do take issue with those who claim that, much more broadly, there is no trade-off. In particular, it has been pointed out that some countries, by "locking down", have been able to sharply reduce deaths and, apparently, avoided loss of output compared to some other countries that applied less strict measures. We believe these are exceptions and reflect particular circumstances that will certainly not hold for all epidemics. Moreover, although economic activity may not have fallen more in some economies with stricter measures, we believe that a broader welfare measure would reveal larger losses. By "welfare", we thus have in mind a concept that not only includes traditional economic variables like job security or consumption but also the value of leisure and life enjoyment more broadly.[2] We must also take account of lost human capital accumulation (which will typically not be included in GDP) in cases where learning centers (such as schools and

[1] This quick containment of course also leads to some losses in welfare but, arguably, they are minor and can be abstracted from.

[2] Here, one would obviously include possible effects on the prevalence of depression, domestic violence, and so on.

universities) are prevented from operating fully. Finally, let us point out that, as is always the case under policy evaluations, the trade-off between costs and benefits arise at the efficiency frontier, i.e, when policies are well designed. In a situation where bad policies are in place, it is possible to improve the outcome along all dimensions.

The second, and related, presumption is that, for the purpose of systematic policy comparison, it is meaningful to use a monetary measure for the "value of life". The point here is not a philosophical one, nor is it that we propose a specific value; rather, we regard it as a tool for systematic evaluation and, most importantly, for ensuring that different policy statements be internally consistent: a high value of life in one context, we posit, should translate into a high value in other contexts as well.

We thus propose an *integrated assessment model (IAM)*: a two-way feedback model with explicit economic activities (including social interaction) alongside the epidemiological spread. We borrow the formulation of economic activity from standard economic theory, just like we borrow the spread of the virus from epidemiological theory (e.g., a SIR model). The two naturally connect provided that the economic model involves a description of social activity and provided the virus affects economic activity (and human welfare) explicitly. Hence, we use the label *epi-econ* to describe our IAM. Another important tenet is that our model be quantitative: it must be parameterized so as to replicate the key interactions in focus. This is a challenging requirement both for the epi and the econ parts of the IAM, but nevertheless a central one.

As perhaps has already become clear, a key feature of the kind of economic theory we have in mind is that is has an element of sociology, i.e., an aspect of economic activity that describes how humans interact at work and while enjoying leisure. Thus, a more appropriate attribute than epi-econ might be epi-socio-econ, but we stop short of including the "socio" term mostly because we may be using it inappropriately. We are also not aware of any quantitative sociological modeling that can be imported and used here.

The model framework we outline thus involves decisions about production, consumption, and leisure. It also specifies explicitly how humans value being alive. The framework features population heterogeneity both with respect to epidemic risk and economic productivity. Since many infections disproportionately hit specific subgroups of the population, and since subgroups of the population vary in productivity, this heterogeneity matters for the quantitative assessment of the effects of an epidemic. From the epidemic side, an epidemic spreads when people meet and socialize, either in the workplace or in their spare time, i.e., not when they consume goods in general. Along the same lines, to evaluate the welfare implications of different actions aimed at restricting the spread of the virus, it is important to take into account the lost welfare from not being able to enjoy leisure in a social context.

In the text that follows, we begin by discussing how an IAM would, in principle, work in Section 2. In Section 3 we list all the features that we think a model that is to be useful for policymakers ought to have. In the following Section 4 we discuss what data would be important input for the model, for its parameterization as well as for its testing. We also include a section with actual data, in this case from the United States, illustrating what directions we need data collections to take, as well as some challenges. In Section 5 we then discuss how an epi-econ model can/should be used, once constructed and tested. We finally present an explicit epi-econ model in Section 6: one we are developing (see Boppart et al. (2020)). This model is merely a concrete illustration of the kind of setting we have in mind: a proof of concept. We also include some results from policy analysis using this model. Section 7 concludes.

2 Integrated assessment

Integrated assessment models of epidemics and economics allow us to perform *counterfactual policy analysis*. I.e., the model is used as a laboratory in which different policies can be tried and outcomes can be recorded and compared. Some laboratory-style studies can also be carried out in practice (e.g., by randomizing treatments) but they are not common and the resulting insights are not necessarily generalizable beyond the specific samples studied.[3] The reliance on a model, of course, is hazardous to the extent that the model is too stylized a depiction of reality. In economics, simple models are in fact often used for the purpose of ensuring logical coherence and for gaining intuition about complex mechanisms. Here, however, the aim is quantitative, and hence elegance and intuition may have to play second fiddle to the demands of a complex reality. Most of the integrated assessment models that have been constructed so far in this new area are, in our view, not yet ready for full quantitative analysis and hence, not yet sound grounds for policy advice. For more on this, see Section 3.3 below.

To understand the workings of an IAM—in particular, the two-way feedback between epidemiology and economics—we first illustrate with a more well-known case: that where economics is analyzed jointly with the climate. The connections to the case of epidemics will then be apparent.

[3]Such studies are also rarely informative of the effects of large-scale implementation: they do not speak to "general-equilibrium effects", i.e., economy-wide repercussions.

2.1 The case of climate change

In the early 1990s, William Nordhaus developed the DICE model: a Dynamic Integrated model of the Climate and the Economy. This was the first full IAM in this important area, and many similar climate-economy models have been constructed since. The fundamental feature of this model is that the emission of carbon dioxide, a byproduct of any economy activity where fossil fuel is used, into the atmosphere causes global warming, because carbon dioxide is a greenhouse gas. Warming, in turn, affects people's welfare through many channels economists study, including consumption, health, and value of their spare time (leisure). Thus, any economic policy aiming to limit the use of fossil fuel works its way through the natural-science part of the IAM and, through contained warming, back to the economy in the form of diminished effects on human welfare.

Figure 1 illustrates. The right-hand side Economy block describes people—how their welfare is determined—and their economic activities, which includes the burning of fossil fuels. Government policy, such as a carbon tax or regulatory limits on emissions (policy 1 in the figure), can be used here to affect the carbon dioxide emissions. These emissions feed into the left-hand side Nature box. This box contains the carbon cycle, describing the process by which carbon dioxide enters, circulates, and exits from the atmosphere over time, and the climate model, which contains the details of the greenhouse effect and warming. Finally, warming, through its effects on humans, feed back on the Economy block. The circular feedback also has a time dimension not depicted here; for example, the carbon cycle plays out over many hundred years. There is, in addition, a spatial dimension: although emissions have the same effect on global warming no matter where they take place, warming, and damages, differ greatly across space. Finally, note that we have inserted policy boxes 2, 3, and 4 as well: Box 2 refers to carbon capture (i.e., the prevention of carbon dioxide from entering the atmosphere, which also requires storage), box 3 to various forms of geoengineering (e.g., ejecting aerosols into the atmosphere or the Fuglesang-Hassler idea of using giant parasols in space to prevent solar radiation to reach Earth), and box 4 to adaptation (air conditioning, the construction of walls to prevent flooding, etc.).

The IAM should thus specify a clear structure allowing us to see how policies act through the entire system. For example, it is common to analyze climate change taking emission paths (over time) as given, but that neglects the feedback that the implied warming has on emissions themselves. In addition, an explicit structure makes apparent what is missing and/or controversial, which leads to further modeling and a scientific process that, in the end, will be crucial for policymakers. The fundamental

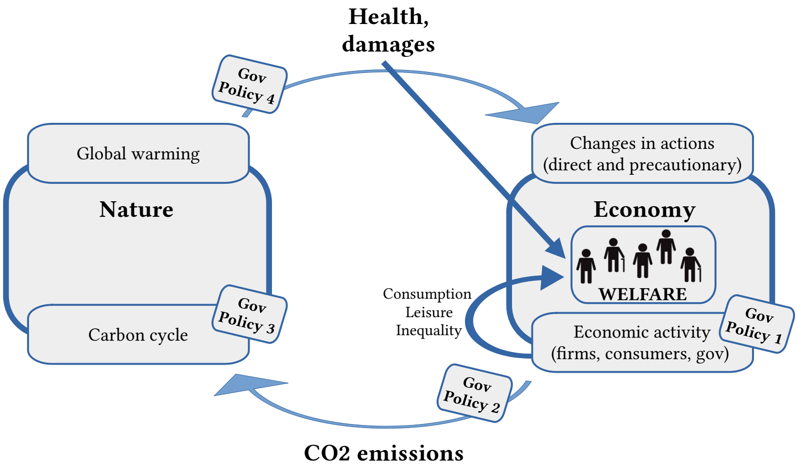

Figure 1: The climate-economy IAM

point here is that whereas it is clear from the general logic that we need to lower carbon emissions, it is far from clear how to best accomplish it. Fossil fuel needs to be phased out, which will occur through a process where markets and political decisions interact. To understand this interaction and device efficient policies, economics is needed.

2.2 Epi-econ IAMs

Like in the case of climate, when trying to understand epidemics there is also key interplay with the economy. Of course, when people stay home rather than go to work or go shopping (because they are told to or merely choose to), production and measured GDP falls. But the connection with the economy, from an economist's perspective, goes much further: although it is perhaps not known outside of economics, it is standard for economists to interpret the term "economy" in a broad sense. First, it includes activities valuable to humans that are not counted in GDP (such as leisure) and there is also a quality component to such activities (leisure is not merely the time spent not working; a proper vacation is typically preferred over a "staycation"). Second, economists also record the evolution of hard-to-measure and yet crucial variables such as the stock of human capital (which increases with schooling as well as learning in market activities) or health. Third, economists do take into account the value of life in analyses, as well as health. The model and data collection we propose here uses economics in this broader sense.

Key in the spreading of epidemics is the meeting of people and, therefore, social interaction. Here, it is important to note that qualitative as well as quantitative aspects of the social interaction are important for the transmission of an epidemic. Social interaction can occur in virtual reality as well as in real life and also physical encounters can be more or less contagious. These different types of interaction are substitutes but not perfect substitutes. One should therefore expect that policies that change qualitative features of social interaction will have (distortionary) costs that need to be weighed against their benefits. From our perspective, it would have been perfect if we could simply borrow a quantitative, off-the-shelf model of such interactions from sociology and merge it with our economic model. Such a model would list interactions and their nature and describe their determinants, and it would include sufficient heterogeneity (such as age, which is key in this context). However, we have not been able to identify any useful candidates—and it may be that they do not exist.[4]

Figure 2 shows how the epi-econ model works schematically. Many items are identical to those in Figure 1 but the Nature box is replaced by a box we label Epidemic—also something from the natural sciences—which contains the core epidemiological theory at play. The key link from the economy to the epidemic is the social interaction that generates spreading and the key link back is the effect on health and lost lives, though of course there are details left out from the figure in these relationships. A key feature inside the Economy box that we will focus on here is the changes in actions and the fact that they have different determinants, some of which are "precautionary" (i.e., once people are aware of the virus, they avoid social interaction even absent recommendations or restrictions).

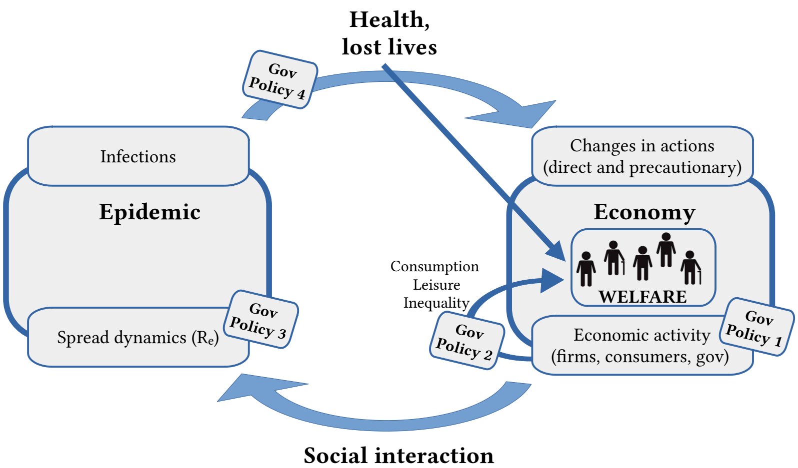

Figure 2: The epi-econ IAM

[4]We are very grateful for any pointers in this direction!

We also want to point out that whereas the epidemiological block is meant to be rather standard, but of course have different specific features depending on the kind of virus in question, the economic one is less standard in one respect: as will be more clear below, we advocate explicit descriptions of social activities, e.g., drawing distinctions between leisure activities that are socially intense and those that are not and between consumption choices that are more or less socially interactive. Although such a distinction is not standard in economic modeling, it poses no conceptual complications. It does, however, require that we specify how different types of time use are valued by individuals and how productive they are in different economics activities. In particular, we will display some data below that informs our model parameters.

Turning to policy, we have introduced four boxes in the picture; there is no perfect way to divide policies and they have multiple dimensions, e.g., some being purely medical and others purely of an economic or sociological character. Let us thus discuss these all in turn (the list is not exhaustive but contains many policies that have been used around the world).

1. Measures with a direct effect on the suppliers of goods and services:
    - restrictions on medium- and long-distance travel
    - recommendations to work at home/remotely
    - the closing of schools and daycare centers and restrictions amounting to remote schooling
    - the closing of restaurants or restrictions on their operations and
    - restrictions on the use/availability of public transportation.

    Note also that some of these can be obtained by incentive devices, such as taxes/subsidies, whereas other require direct restrictions.

    In this category, we also find all those government interventions aimed at supporting households and firms financially until the epidemic is over: so-called bridging policies. These are thus caused by the epidemic but purely economic; most of these policies aim at reducing the negative economic side effects of policies and in the voluntary changes in behavior to reduce contagion. Here the idea is not to make consumers demand more, as in a normal recession caused by shortfalls in demand, but rather to help consumers maintain a decent standard of living despite an income shortfall. Similarly, it is not about "raising restaurant demand" back up again but to support restaurant owners financially in order to avoid bankruptcies that would make rapid recovery after the pandemic more difficult.
2. Measures restricting leisure activities involve

- travel restrictions (as above)
- limits on the number of people congregating in groups
- special recommendations for retirees (possibly specifying by age)
- restrictions on visitation at nursing homes
- limits on visiting restaurants and other service establishments, cultural events, etc.
- public transportation (as above) and
- general recommendations to keep distances.

Here, total (or almost-total) lock-down amounts to the simultanteous use of most of the categories under 1 and 2.

3. Non-pharmaceutical interventions with negligible direct economic or welfare effects. These include
   - recommendations on hygiene
   - recommendations to wear face masks/shields (belongs in category 2 if non-negligible utility costs are incurred) and
   - generally maintaining social distances.
4. Measures to mitigate the consequences of infection include
   - better treatment (which is knowledge-based and require the purchase of potentially expensive drugs) and
   - investment in physical capital that increases the capacity to treat seriously ill patients (intensive-care units).

As pointed out, some of these items can be achieved, in part at least, with economic incentives. In Sweden, for example, sick leave is currently granted with full pay from the first day—in contrast to normal times, when it applies only from day two and on.

Finally, a policy that perhaps belongs in group 3 and that has been used broadly around the world, though less systematically in Sweden, is “test, trace, and isolate”. The idea here is targeted (or potentially even broad) testing and then tracing, with isolation of those testing positive or having been in contact with those testing positive.

Not all of the policies can be studied in the framework we have developed ourselves—because it is still a stylized model—but a satisfactory model for policy evaluation should allow most of them to be evaluated quantitatively. In the next section, we describe in somewhat more detail what elements are necessary in an epi-econ IAM, while still remaining at a conceptual level. Our own formal mathematical model is then presented in Section 6 as a concrete example of the kind of structure we call for.

3 What to aim for in an epi-econ IAM: theory

Let us first describe the epidemiological setting we envision and then move to the economic model.

3.1 The epidemiology model

The very simplest epidemiology model is the SIR model. It can be traced back to Hamer, Ross (early 20th century) and Kermack and McKendrick (1927) and is described compactly in Atkeson (2020) a first application to economics can be found in Eichenbaum, Rebelo, and Trabandt (2020a). It keeps track of three parts of the overall population at each point in time: S, the susceptible, not yet ill people, I, the infectious people, and R, the people who are recovered (and are no longer neither susceptible nor infectious). Also, some people die and are removed from the population. Thus, $S+I+R$ is equal to the total population N, which we can think of as the whole population alive.

Over time, these variables evolve. We use a discrete-time setting in our illustrations and hence specify S_t, I_t, and R_t at time t, where we think of t measured in days from the onset of the epidemic. The simplest SIR model then specifies S_{t+1}, I_{t+1}, and R_{t+1} mechanically as a function of S_t, I_t, and R_t: a certain, exogenous fraction of the S_t group moves into the I_{t+1} group and an exogenous fraction of the I_t group moves into the R_{t+1} group.

The first of these two transitions is more nontrivial: how large the $S_t \rightarrow I_{t+1}$ transition is depends on the size of the I_t group: it stems from the (social) *interaction* of these groups. I.e., we define the *transmissions* T_t between S and I by

$$
\begin{aligned}
S_{t+1} &= S_t - T_t && (1)\\
T_t &= \pi I_t \frac{S_t}{S_t + I_t + R_t}, && (2)
\end{aligned}
$$

with I in turn also affected by the transitions out of this group:

$$
\begin{aligned}
I_{t+1} &= I_t(1-\pi_r-\pi_d)+T_t && (3)\\
R_{t+1} &= R_t + \pi_r I_t. && (4)
\end{aligned}
$$

Here, a basic parameter is π, describing the spread between S and I. Specifically, π is the product of two factors. The first is the number of contacts an infected individual has with other people per period. The second is the probability that this leads to a transmission of the disease if the contact is with a susceptible individual. Thus, πI_t multiplied by the share of susceptible in the population, (S_t/N_t) is the number of newly infected per period. Two other key parameters are π_r, describing the rate at which infected individuals people recover, and π_d, the rate at which they die.

A commonly mentioned parameter is R_0. The notation here is unfortunate; it is not the number of recovered at time zero but something entirely different: it is the so-called basic reproduction number. This equals the expected number of secondary infections generated by a single, typical infection in a completely susceptible population, i.e., right in the beginning of an epidemic. With the notation in the SIR model above, R_0 is given by π/π_r.[5] Why is that? The duration of the infection for the average person is given by $1/\pi_r$. Over one unit of time, an infected individual has π contacts close enough to result in a transmission of the disease. A fraction $S/(S+I+R)$ of these contacts is with susceptible individuals, which in a completely susceptible population is equal to 1. Thus, a single typical infection in a completely susceptible population spreads the disease to π/π_r individuals.

If the basic reproduction number is lower than 1, there will be no epidemic. The logic is straightforward: if an infected individual creates less than one new infection, the epidemic will die out by itself.[6] Figure 3 below illustrates. We assume an exogenous influx of 0.1% infected into the population at time $t = 0$ (travellers from Italy, for instance). In Subfigure 3a, we plot the number of susceptible, infected, and recovered in the population over time if $R_0 = 2.0$. These three groups sum to 1.0 at all points in time. As the figure shows, the number of infected grows exponentially in the beginning, until the population reaches the herd immunity threshold, and thereafter the number of infected starts declining. The herd immunity threshold is reached when the share of susceptible, $S_t/(S_t + I_t + R_t)$, reaches $1/R_0$. To see this, insert Equation (2) for the number of transmissions T_t into Equation (3) using the approximation that $\pi_d = 0$. The number of infected is increasing, i.e., $I_{t+1} > I_t$, if and only if $S_t/(S_t + I_t + R_t) > \pi_r/\pi = 1/R_0$. We can verify in the figure that this value is roughly 0.5.

Furthermore, since the number of susceptible people falls over time, the number of individuals every infected transmits the disease to also falls. Eventually, the pandemic thus dies out because of a lack of susceptible individuals. This occurs before everyone has become infected. Denote the share of people who is expected to remain uninfected

[5] We ignore deaths here, i.e., $\pi_d = 0$ for simplicity; this is not problematic quantitatively, since π_d is very small compared to π_r. Including it would mean that $R_0 = \frac{\pi}{\pi_r + \pi_d}$.

[6] Formally, these statements should be stated in terms of probabilities. I.e., a given agent could end up infecting many people, but such an occurrence would be rare if $R_0 < 1$; the key point is that, in a large population, on average people will infect less than one person when $R_0 < 1$.

forever by x. It can then be shown that x satisfies $x = \exp(R_0(1-x))$. The solution falls in R_0. For R_0 equal to 1.5, 2, and 3, x will be 0.417, 0.203, and 0.060, respectively.

One of the key insights from the SIR model is that the share of individuals who do not become infected, x, is smaller than the herd immunity threshold $1/R_0$. The number of infected begins to decline when the share of susceptible individuals is below $1/R_0$ and if somehow infections were brought down to zero, the epidemic would not resurface after herd immunity has been reached. However, when the number of susceptible individuals passes the herd immunity threshold, there are still many infected individuals in the population and they will infect more individuals than necessary to reach herd immunity. The dynamics of the SIR model therefore features "overshooting": The momentum of the epidemic leads to more individuals becoming infected than necessary to reach herd immunity.

Subfigure 3b shows the same SIR model, but with $R_0 = 0.9$. As can be seen, in this case there is no epidemic: the fraction of susceptible stays at approximately one, and the epidemic does not take off. Subfigure 3c zooms in to show the evolution of infected and recovered. The number of recovered is still increasing as time goes by: remember, each infected still infects 0.9 new persons (approximately, since the population is not completely susceptible). However, since $0.9 < 1.0$, the number of infected is decreasing even in the beginning.

The above description hopefully makes clear that one can summarize the dynamic, day-to-day evolution of the virus with an extension of the model outlined above. Clearly, however, the simple SIR model lacks a large number of realistic features. For one, we need to distinguish heterogeneity across people in several dimensions. One dimension is the propensity to be infected, which appears to differ across individuals. Another is the probability of dying conditional on infection both across individuals and over time (there is learning and economic resources that can be put in place to lower π_d). Still another is the infectiousness, i.e., the propensity to pass the virus on to others, which seems to vary greatly; in fact, keeping social interactions fixed, some people appear to be "super-spreaders" whereas others barely spread the virus at all. But people also interact differently. The *extent* of social interaction will be of specific importance below, where we will describe the dynamic model block. It should of course be noted, too, that all these differences have observable determinants such as age, gender, and socioeconomic status.

Also, both the way the virus spreads and how the characteristics of the virus changes over time (both through, e.g., mutations and seasonality) appears to be partly unknown. For example, the extent to which it is airborne is crucial for modeling the

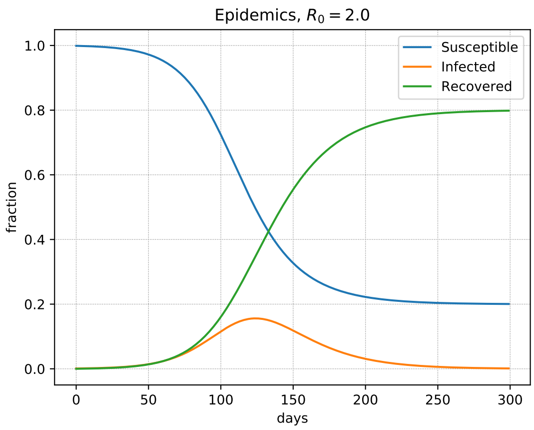

(a) $R_0 = 2.0$.

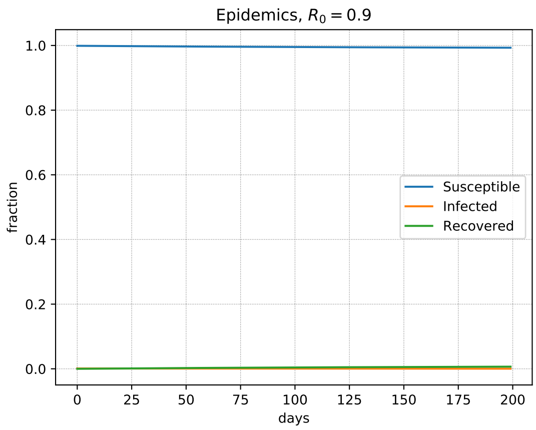

(b) $R_0 = 0.9$.

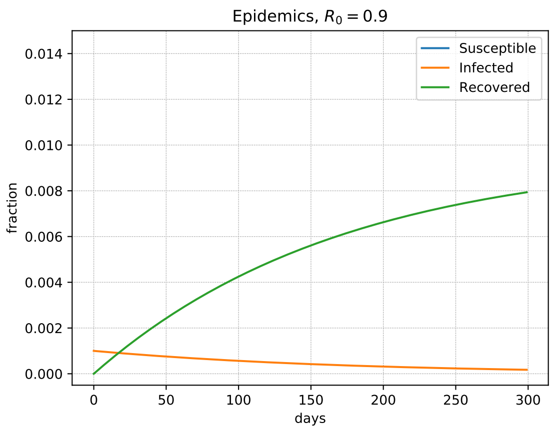

(c) $R_0 = 0.9$. Note scale of y-axis.

Figure 3: The evolution of an epidemic with different R_0.

details of the epidemiological spread. Covid-19, moreover, seems to spread in clusters (in part accounted for by super-spreaders), but it is beyond the scope of the current text to describe exactly how these clusters work.

Much about how the current virus spreads is not fully understood. E.g., it is controversial what the most effective personal protection might be is controversial (face masks and shields, desinfectants, etc.). Relatedly, we do not know how human behavior changes once personal protection is used—does risk-prone behavior rise or fall? The existence of uncertainties is sometimes used as a argument not to formulate IAMs. Our view is that IAMs are still extremely useful. They can be formulated, and policies can evaluated, systematically under different parametric assumptions, providing us with a much better understanding of possible consequences of different options.

3.2 The economic model

The economic framework consists of a set of individuals who are long-lived, relative to the time horizon studied. These individuals act like standard "economic agents": they work and enjoy leisure, and they consume a set of goods and services. They also save, though in the simple model we have outlined below, saving decisions are abstracted from. Consumers' choices are guided by their utility functions and their budget sets, so the baseline approach we advocate is to use standard revealed-preference arguments for normative analysis: we back out, by observing consumers' choices and budgets, which contain market prices, the utility functions behind observed behavior. Then, as we conduct policy analysis, we can make normative statements from the perspective of each individual, since we can evaluate the utility consequences of the policy, once the full model has been specified and solved. In the presence of poorly functioning markets, of course policy might be useful to improve on the utility outcomes, but we will typically base our analysis on a market structure that, apart from the complication implied by the epidemic, works without frictions. Thus, so far, this is all standard economic analysis.

The economic model also has production technologies and firms. These will typically be entirely standard and borrowed from microeconomic theory, and a baseline will be one where there are no frictions in the production sector. In an epidemic where there is a sharp reduction in social activity, and hence also in economic activity, there will likely be unemployment: people who, in principle, are willing to work but cannot work since many activities are (at least temporarily) shut down. There will likely also be bankruptcies and close-downs; hence installed physical capital will not be fully utilized. An economic model clearly ought to have these features of market frictions too (although our simple model below does not). A model without frictions can perhaps be thought of as an indication of what an ideal allocation, towards which policy should aim, would look like.

In an epi-econ model, it is important to connect the consumer's economic behavior to the social activities that underlie the spreading of the virus. First, one would ideally have households in the economic model consisting of more than one member, since a large part of the spread occurs between household members (in the simple model we will describe below, we do not have this feature). Second, one must specify the degree of social activity involved in the various economic choices a consumer makes. In the first epi-econ model, Eichenbaum, Rebelo, and Trabandt (2020a) simply assume that social contact is proportional to aggregate consumption. In the model we outline below, we advocate a more ambitious approach (and we use the simplest possible version of it): to specify that there are different consumption goods, as well as different leisure activities, that differ in their degree of social activity. The idea is that some services, like restaurant visits and vacation trips, are socially active, whereas watching a TV (which involves buying a TV set and paying for a TV or internet license) is much less so. Thus, we envision an economic theory that has a dimension normally absent in economic models, connecting consumer choices (of goods and services as well as leisure) to their degree of social interaction.

In an epi-econ model it is also, as emphasized at the start, crucial to include health variables, or at least deaths, as an outcome as well as something individuals care about. These are standard variables in health economics. There, it is commonplace to incorporate a separate "value of life" in consumers' utility function. I.e., people derive utility from merely being alive, in addition to the value they experience from leisure and consumption. In the simple model we formulate below—and, we think, also in more ambitious models—utility is additive in two terms where the first term summarizes the benefits experienced from consumption of various goods and the enjoyment of various active leisure choices. The second term is then the separate value of life, i.e., the value of being alive during that period of time. Hence, a young consumer dying from covid-19 loses more in utils than an old consumer, as the former loses more years of remaining life-time. One can also introduce health in a way that would affect the value of life (a "quality" component, e.g., making it more valuable to have a year of life with perfect health than with ailments); health can obviously also affect the utility value of consumption and active leisure.

A final component that needs discussion is the information individuals have. Information here has two interpretations. One is the information an individual has about his or her own health (with regard to the virus), and similar information about the health of others. Similarly, it makes a big difference if the government knows who is infected and who is not. Precisely, test, trace, and isolate strategies are a means of acquiring information and acting on it. In our models, these strategies can be formulated, but fully carrying out such analysis is quite difficult as it would entail keeping track of all individuals and their health status over time, which makes most models intractable. However, one can still run simple experiments that are informative on the value of this class of strategies.

Another kind of information involves expectations: do people know what will happen to the virus over time? Of course they will not be able to predict who will be infected and when but do they have accurate statistical expectations on the evolution of the virus in the overall population? Relatedly, do they have an accurate picture of the probability of dying and how it evolves over time? The model needs to take a stand. We illustrate this point with our simple model below where we compare perfect foresight—the case where all fundamental features of the virus is known and hence infection and death probabilities must evolve deterministically over time as a function of the chosen levels of social interaction of all individuals—to one where people behave entirely myopically, i.e., where they act as if they are unaware of the virus, or do not believe it is dangerous for humans. Clearly, here, it matters greatly to outcomes whether a vaccine, or cure, arrives and, if so, when, as we shall also illustrate.

The above discussion lists all components of the desired economic model: in short, a set of (heterogeneous) consumers, their preferences (including a value of life per time period), firms, an accompanying market structure (with some frictions), and a specification of information. We will be very explicit in describing our own simple such setup below.

3.3 Epi-econ modeling: a very brief literature review

In this section, we very briefly review the explosive epi-econ literature emerging after the onset of covid-19. We restrict the review primarily in two ways. First, we do not review the purely empirical literature (documenting patterns of virus spread and economic activity, over time and space). In some ways, this is a problematic omission, since we do argue that collecting these kinds of data and testing models against them is very important. However, very few studies so far connect epi-econ models with the new evidence—we do list some of them below—and we therefore found this limitation appropriate. Second, we do not systematically review new contributions to epidemiological modeling (some of which have been made by economists).

Atkeson (2020) describes the core model developed by epidemiologists (Kermack and McKendrick 1927). Epidemiological models had been used in economics prior to covid-19 in analyses of other viruses (e.g., Geoffard and Philipson (1996), Kremer (1996), Adda (2007), Chan, Hamilton, and Papageorge (2016), and Greenwood, Kircher, Santos, and Tertilt (2019)), but Eichenbaum, Rebelo, and Trabandt (2020; ERT) was as far as we can tell the first application of epidemiology that is also macroeconomic in the sense that it describes a whole economy of forward-looking agents and market determination of prices. ERT build a wholly microeconomic structure (including rational consumers and firms with objectives explicitly described) and the interventions they consider involve fully described policy instruments and comparisons between laissez-faire and fully optimal policy. It is also quantitative in that the model's economic parameters are selected to match (standard) characteristics of macroeconomic data—and the epidemiological parameters are chosen to match known estimates pertaining to the specific features of covid-19.

We organized the review of the literature that emerged after ERT as follows: we used Google Scholar to identify papers that cite ERT, we removed the papers that did not contain epi-econ modeling, and then ordered the papers by citation count. We carefully reviewed the top 30 papers on this list and we then added other papers we have become aware of that, for some reason, did not appear on our list. We are aware that our literature review is highly likely to miss some important contributions but we plan to update our list of papers accordingly when we identify further relevant papers. Next, we decided on five important "features" and categorized the papers accordingly: for each feature, we list which papers have the feature. The result can be found in Table 1.

Feature 1 is quantitative ambition. Many papers in the new literature examine the logic of the interaction between epidemiology and economics, or the logic how one might want to optimize welfare in an epi-econ setting. These papers are interesting and clearly useful for our understanding of how the models work, but for the present purposes we consider quantitative evaluation critical.

Feature 2 regards the information structure: a paper is in this category if it allows people in the model to be aware of their health status (whether they have, or have had, covid-19), potentially also allowing testing. Another issue here is whether the policymaker has this information about individuals. Clearly, this feature is relevant if we want to evaluate the test-trace-isolate policy prescription. However, since the whole collection of health statuses in a population is high-dimensional and its dynamic evolution—especially when it is endogenous due to testing choices—is correspondingly

Feature	Papers
Quantitative ambition	Alon, Baron, Bar-On, Cornfeld, Milo, and Yashiv (2020); Alvarez, Argente, and Lippi (2020); Aum, Lee, and Shin (2020); Bethune and Korinek (2020); Bodenstein, Corsetti, and Guerrieri (2020); Bognanni, Hanley, Kolliner, and Mitman (2020); Brotherhood, Kircher, Santos, and Tertilt (2020); Eichenbaum, Rebelo, and Trabandt (2020a,b); Farboodi, Jarosch, and Shimer (2020); Giagheddu and Papetti (2020); Giannitsarou, Kissler, and Toxvaerd (2020); Glover, Heathcote, Krueger, and Rios-Rull (2020); Jones, Philippon, and Venkateswaran (2020); Kapicka and Rupert (2020); Kaplan, Moll, and Violante (2020); Krueger, Uhlig, and Xie (2020)
Consider a version where people or the planner is informed about the R identity of agents	Acemoglu, Chernozhukov, Werning, and Whinston (2020); Acemoglu, Makhdoumi, Malekian, and Ozdaglar (2020); Alvarez, Argente, and Lippi (2020); Aum, Lee, and Shin (2020); Bethune and Korinek (2020); Brotherhood, Kircher, Santos, and Tertilt (2020); Eichenbaum, Rebelo, and Trabandt (2020b); Farboodi, Jarosch, and Shimer (2020); Kapicka and Rupert (2020); Krueger, Uhlig, and Xie (2020); Piguillem and Shi (2020)
Model founded on a microeconomic structure	Aum, Lee, and Shin (2020); Bethune and Korinek (2020); Bodenstein, Corsetti, and Guerrieri (2020); Brotherhood, Kircher, Santos, and Tertilt (2020); Chang and Velasco (2020); Eichenbaum, Rebelo, and Trabandt (2020a,b); Giagheddu and Papetti (2020); Glover, Heathcote, Krueger, and Rios-Rull (2020); Jones, Philippon, and Venkateswaran (2020); Kapicka and Rupert (2020); Kaplan, Moll, and Violante (2020); Krueger, Uhlig, and Xie (2020); Moser and Yared (2020); van Vlokhoven (2020)
Policy evaluation, comparison between laissez-faire and optimum (potentially with a constrained planner)	Acemoglu, Chernozhukov, Werning, and Whinston (2020); Acemoglu, Makhdoumi, Malekian, and Ozdaglar (2020); Alon, Baron, Bar-On, Cornfeld, Milo, and Yashiv (2020); Alvarez, Argente, and Lippi (2020); Aum, Lee, and Shin (2020); Bethune and Korinek (2020); Bodenstein, Corsetti, and Guerrieri (2020); Brotherhood, Kircher, Santos, and Tertilt (2020); Chang and Velasco (2020); Eichenbaum, Rebelo, and Trabandt (2020a,b); Farboodi, Jarosch, and Shimer (2020); Garibaldi, Moen, and Pissarides (2020); Giagheddu and Papetti (2020); Giannitsarou, Kissler, and Toxvaerd (2020); Glover, Heathcote, Krueger, and Rios-Rull (2020); Jones, Philippon, and Venkateswaran (2020); Kapicka and Rupert (2020); Kaplan, Moll, and Violante (2020); Krueger, Uhlig, and Xie (2020); van Vlokhoven (2020)
Testing the model against real-time data	Aum, Lee, and Shin (2020); Bognanni, Hanley, Kolliner, and Mitman (2020); Farboodi, Jarosch, and Shimer (2020); Giagheddu and Papetti (2020); Krueger, Uhlig, and Xie (2020)

Table 1: Overview of key epi-econ papers. See text for more information about the selection process of the papers and a description of the features.

complex, it is very demanding to analyze models that fully admit this feature. Thus, all the papers we list in this category do not have a full treatment of this important information feature, but they do contain some element of it.

Feature 3 is straightforward: it just lists the papers that have explicit microeconomic structures, thus specifying agents' utility functions, and so on. In contrast, some papers have stylized elements, such as ad hoc rules for behavior or explicit maximization given a objective function that is not described in economic terms. These papers are not excluded from this category because they do not offer valuable insights, but because they use models that are not micro founded. Relatedly, Feature 4 are papers that compare laissez-faire to an optimum, or an outcome with some form of policy instrument. Many papers instead focus on one or the other and, clearly, there are good reasons for this; more detail can be allowed in a positive model if normative analysis is not required (this is especially the case in the applications of private information about health status). Similarly, sophisticated and highly relevant policy options can be developed and studied in detail but only if the economic model is not too complex. Feature 5, finally, contains papers that do make contact with the recent high-frequency data and use it to either test their models or restrict model parameters.

More features could, obviously, be identified. We have argued in this paper that population heterogeneity is key: that it is important to identify risk groups (e.g., young vs. old) and perhaps also other characteristics. Relatively few papers have this feature.[7] We have also argued that it is key to make the trade-off between the welfare of those alive and being alive by itself explicit. To us, surprisingly few papers include an explicit account of the value of life.[8] We also think it is important to consider further epidemiology- or medicine-related features, e.g., a careful discussion of when a cure/vaccine arrives, how death rates depend on the pressures of the medical system, the existence of super-spreaders or cluster effects, the importance of seasons in affecting the spread, and so on. We did categorize papers in these dimensions too but do not report the results here; they are, of course, available upon request. A set of papers also have focus on inequality and, consequently, on possible related policy options (such as increased social insurance). We have not been able to distinguish epi-econ papers that look at richer macroeconomic models with credit and finance frictions (hence allowing explicit analysis of so-called bridging policies); however, we suspect that for

[7]They include, among others, Acemoglu et al. (2020), Brotherhood et al. (2020), Giagheddu and Papetti (2020), and Aum et al. (2020). van Vlokhoven (2020) includes population heterogeneity along the social dimension: some people are more social and interact with more people.

[8]Those that do discuss the value of life include, among others, Alvarez et al. (2020), Alon et al. (2020), and Krueger et al. (2020). A discussion can also be found in Hall et al. (2020).

these purposes the "epi" part may not be crucial.

4 What to aim for in an epi-econ IAM: data

The construction of a model is worth it for several reasons. One is to simply try to account for what we observe, or to try to predict the future. A second is for studying counterfactual policy experiments. A third is for discussing what policy might be optimal (with a specific social welfare function in mind). All of these tasks are quantitative in nature so it is important to critically evaluate the modeling choices as well as the specific values for the parameters (of the epidemiology and economics structure assumed). To do this, we must collect reliable data to be used to inform the model's parameter values and test its predictions. Some of these data are readily available but much is not.

First, tracking an epidemic involves data collection at a high frequency—perhaps weekly or even daily data. It is thus important to make every effort to do this. During the covid-19 pandemic, such data collection has taken place around the world, in Sweden and elsewhere, and we must now make maximal use of it and collect further data if necessary.

Second, we will remain silent here on purely epidemiological data (e.g., medical data), given that it is not our expertise. But, needless to say, such data is critical for modeling the epidemiological part of our setting.

Third, as for economic variables, what is most challenging is arguably linking up consumer choices to social activities—and the extent to which they involve close interactions. For this reason, we include a rather detailed discussion of such data below from the United States: time use surveys. Similar data is available for Sweden and we cover the case of the U.S. here because it is be used for assigning parameter values in our simple model below.

Fourth, when it comes to the "testing" of the model, we would like our theory to accurately describe how much the production of goods and services (aggregate output) will fall if various restrictions are adopted and how different people change their time allocations across activities.[9] Clearly, for this, it is very important to use high-frequency data, ideally in the form of natural experiments, if available. Absent this, one can still—again, as will be illustrated for our simple model—assess what seems a priori reasonable.

[9] We are aware that the National Institute of Economic Research has begun looking into this.

4.1 Time use data: the United States

In this section, we illustrate how time-use data can be put to productive use in this area. We describe some available U.S. data, particularly with an eye toward our simple epi-econ IAM, to be described below. In this model, we make a distinction between people of different ages (given that covid-19 vulnerability has an important age dimension) and we take the coarsest possible perspective: we distinguish "young" from "old". Thus we define these two groups and look at how they spend their time on leisure vs. work, and, within activity, how much of the time is spent in a socially active way. We focus on the economy before any pandemic hits, in order to understand the pre-pandemic equilibrium allocations. However, our classification of activities is guided by the degree of social interaction, which is of key importance once the pandemic hits.

We use the American Time Use Survey (ATUS), which provides nationally representative estimates of how and where Americans spend their time. Importantly, it includes data on the full range of nonmarket activities, from relaxing at home to restaurant visits and attending sports events.

We want to include the full population in our model, and not only for example the working-aged. The reason is obvious: the epidemic does not only affect the working-aged, quite the opposite. For our purposes, we define the young as individuals aged 15-60, while the old are those above the age of 60.[10] This classification is mainly guided by epidemiological concerns: the epidemic is shown to hit the older population much harder, and there 60 seems to be a reasonable cut-off. The reason for not including the population below 15 is twofold: first, currently, it does not seem that the epidemic is driven by children in an important way. Second, the ATUS does not sample individuals below 15, and we do not think we add information to our analysis by trying to impute children's behavior.

4.1.1 Leisure vs. work for young and old

First, we look at the amount of time spent on work vs. leisure. We divide the 24 hours a day in three mutually exclusive and complementary exhaustive broad categories: sleep, work, and leisure.

Sleep is defined as the time spent either sleeping or experiencing sleeplessness.

Work we define as the sum of the following activities: market work, core housework (meal preparation and cleanup, doing laundry, ironing, dusting, cleaning, etc.), other home production (home maintenance, outdoor cleaning, vehicle repair etc.), necessity

[10] This was a painful decision, given that two of the authors of this paper have reached 60.

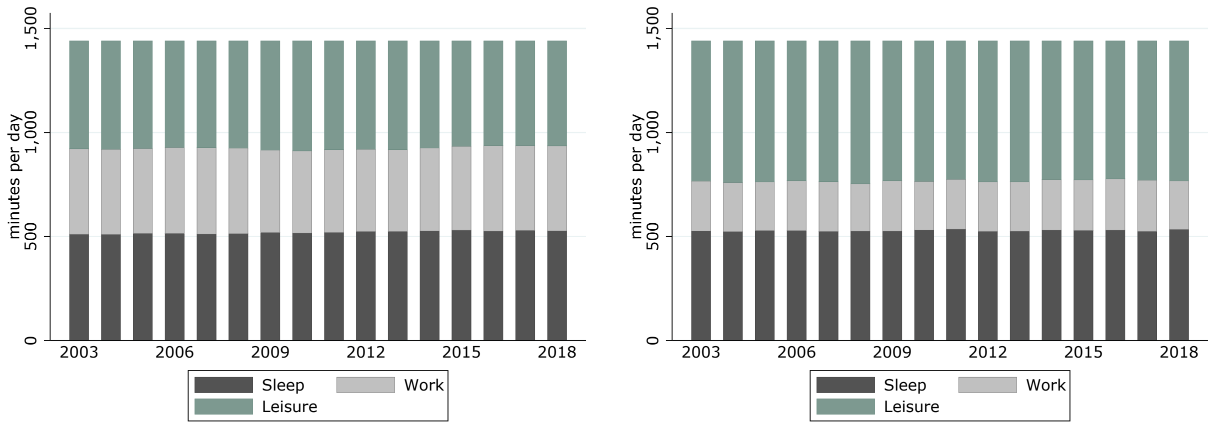

(a) Young

(b) Old

Figure 4: Time spent on sleep, work, and leisure over time. A full day is $24 \cdot 60 = 1440$ minutes.

shopping (grocery shopping, going to the bank, etc.), and time spent in education. We also add all travel time associated with any of those activities.

Leisure, lastly, is defined as the sum of the following activities: entertainment/social activities/relaxing, child care and caring for other adults, gardening, time spent with pet, personal care, eating and drinking, recreational shopping, civic and religious activities, and own medical care. Again, all travel time associated with any of those activities is added to the total.[11]

Figure 4 shows how a day is divided into the three categories sleep, work, and leisure for young vs. old between the years 2003 and 2018. A first observations is that there are no big movements in the time use during this period. A second observation is that young and old spend very similar amount of time sleeping.[12] A third observation is that, not surprisingly, old spend less time working than the young. In our definition of work we also include housework and home production, so time spent on work is not negligible for the old, but still the difference is large.

[11]This definition of leisure is close to leisure Measure 4 used by Aguiar and Hurst (2008). Compared to that definition, our leisure concept adds recreational shopping, gardening and time spent with pet, but excludes sleeping and education. In the category leisure shopping we include "Shopping, except groceries, food, and gas", "Comparison shopping", and "Researching purchases, n.e.c.".

[12]Old individuals sleep on average more than individuals in the core working age (30-60), but our definition of young includes also the very young, who are heavy sleepers.

4.1.2 Leisure spent at home vs. outside home

Next, we want to understand how much of the leisure time is spent in socially intensive activities. For our purposes, we define socially intensive activities as activities spent outside the home. We prefer this classification to the alternative "with whom" criterion, since we consider, e.g., the activity of going to the mall for recreational shopping to be a socially intense activity, even though the individual goes there on his/her own.

We classify activities as *not* socially intense if it took place in the respondent's home or yard. Moreover, we classify personal care activities (e.g., grooming and personal activities) coded with location code "Blank" in the survey as not socially intense. Lastly, 0.3% of the observations in the data are coded with "Unspecified place". For these observations, we code those where it is plausible that the activity took place in the home as not socially intense.[13] The socially intense activities are consequently the activities that took place outside home. Examples of locations for these activities include someone else's home, store/mall, restaurant or bar, and gym/health club.

Figure 5 shows the time spent in socially intense leisure vs. not socially intense leisure for young and for old. As can be seen, despite spending much more time on leisure in total, the old spend approximately the same amount of time on socially intense leisure, i.e., leisure outside their home, as the young.

To understand what these broad categorizations mean in practice, Figure 6 shows socially intense leisure and not-socially-intense leisure broken down on a finer level. For instance, the category "Eating and drinking" shows up in both types of leisure: young spend on average 34 minutes per day eating and drinking outside their home (socially intense leisure), and 35 minutes on eating and drinking at home (not-socially-intense leisure). The largest category for leisure is "Socializing, relaxing, and leisure", both when it comes to socially intense leisure and not-socially-intense leisure. On a finer classification level, the most common subcategory within "Socializing, relaxing, and leisure" for the socially intense type is "Socializing and communicating", while it for the not socially intense type is "Relaxing and leisure", which roughly translates to watching television at home.

4.1.3 Working time spent at home vs. in the office

We now turn to where work time is spent. As described above, our definition of work includes market work, household work, core housework and home production,

[13] As an example, we code the activity "Caring for and helping household children" as not socially intense, while "Participating in sports" is classified as socially intense.

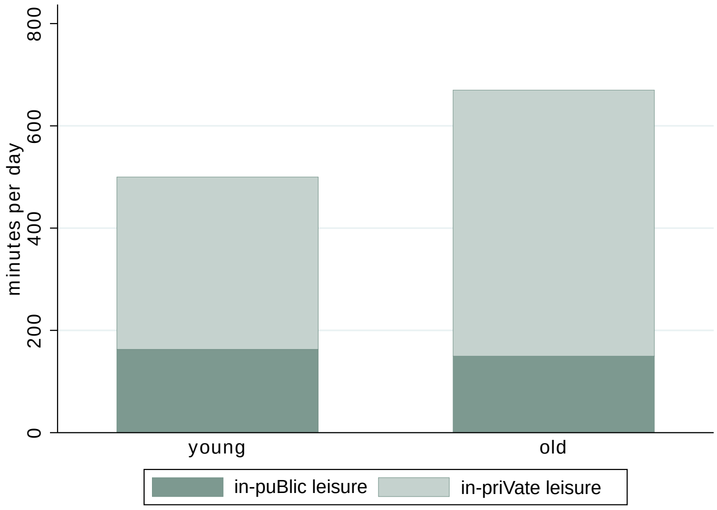

Figure 5: Average minutes per day spent in socially intense leisure activities and not socially intense activities (includes associated traveling). A full day is $24 \cdot 60 = 1440$ minutes. Source: ATUS 2018.

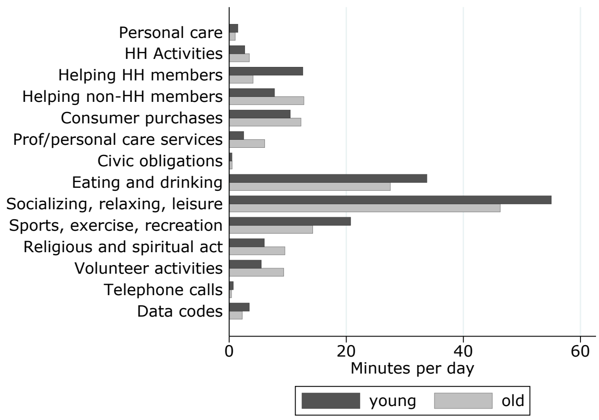

(a) Socially intense leisure

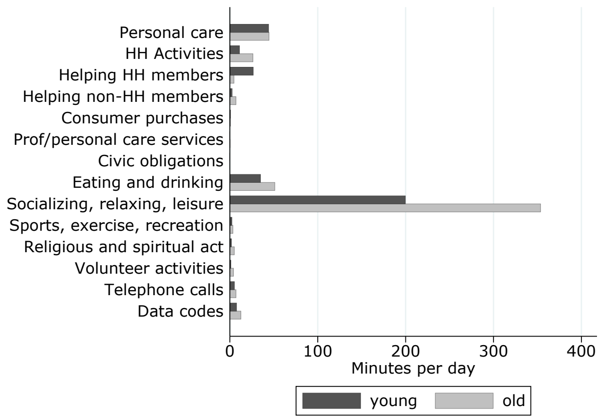

(b) Not socially intense leisure

Figure 6: Average minutes per day spent in socially intense leisure activities and not socially intense activities, by two-digit categories (includes associated traveling), year 2018. "Data codes" refer to observations where the respondent couldn't remember or refused to answer. A full day is $24 \cdot 60 = 1440$ minutes.

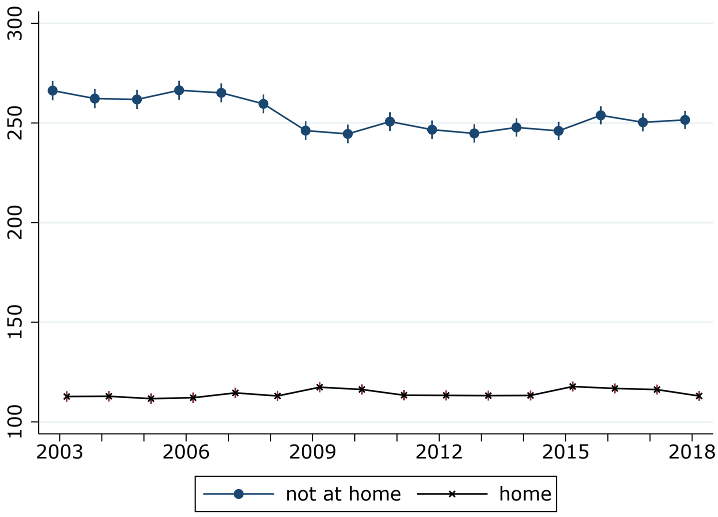

Figure 7: Average minutes per day spent working home vs. not from home. Source: ATUS.

necessity shopping, and time spent in education. In the same way as for leisure, we classify all work activities according to where they were performed: in the home or outside home. Figure 7 shows how many minutes of the average working day is spent at home (112 minutes on average) and outside home, mainly at the workplace (255 minutes on average). Note that some market work is done from home, and some of the household work is done outside home (for example shopping necessities for the household). The same is true for education: some activities are done outside home (mainly in the school) while others (for instance homework) are done at home even in pre-pandemic times. As Figure 7 also shows, there is surprisingly little time trend in how large fraction of the working time that is spent at home. The slight downward trend in work done outside home can mainly be attributed to a compositional effect: the fraction of old individuals has increased slightly during this time period (21% in 2003 compared to 27% in 2018), and they work less, especially outside the home.

4.1.4 The nature of employment

Finally, we turn to the question how the workforce is divided between producing goods and services that involve socially intense interaction vs. production that does not. We use employment statistics on the 4-digit NAICS level from BLS to classify sectors in the U.S. In our simple model below, we again classify jobs in two groups by "sector": those producing goods or services, like restaurant meals, that involve intense social interaction (both in terms of the leisure time and the restaurant time itself) and those that do not. The classification below is thus based on whether the sector is assumed to provide goods/services to the socially intense consumption-leisure bundle (and consequently if the workforce interact with customers). The extent to which the sector can be classified as socially intense can be fully (100%), to a high extent (75%), to a somewhat smaller extent (50%) or not at all (0%).

We then sum up the affected workforce, and obtain that out of the total workforce (161,037,700 workers), 20% work with producing for the socially active bundle.

The lion's share (43%) of the workforce working in the socially intense sector is working in the accommodation and food services provision, followed by "all other retail" (18%) and "non-agricultural self-employed" (14%).

5 How to solve and use the model

Having specified a full IAM and assigned numerical values to its parameter values (we discuss parameter choice below in the context of our example), the model of course needs to be solved. This must be carried out on a computer. The more complex the model is, the more demanding will the task of solving it be (in terms of programming as well as running time).

A key feature of the framework we propose also complicates this task: forward-looking beliefs. I.e., an epi-econ model plays out over time and forward-looking beliefs are crucial, since people worry about risks in the future.[14] In concrete terms, an epidemiological model (and any other pure natural-science processes, like climate change) can be fully solved forward: there is a state of nature at time t (in the SIR model the state is the values of these three variables) and one can then mechanically figure out what happens at time $t+1$, without knowledge of the future. In economic models with forward-looking, people's decisions at t instead also depend on what they expect will happen after t, which includes the behavior in the future: the behavior in the future thus feeds back on the present. Thus one is forced to use solution techniques that *jointly* solve for behavior at all points in time. Decades of work in macroeconomics have fortunately given us methods for these cases, and these methods are now applied and tailored to the new epi-econ field.

Turning to the use of the model, one can first simply try to make forecasts based on "laissez-faire" scenarios (in the climate IAMs, the term "business as usual" is stan-

[14]Whether they are fully rational in this regard or not is not critical here: any amount of forward-looking gives rise to this challenge.

dard). Forecasts are obviously useful for many purposes. In economics, short-run forecasts tend to be better the more information and complexity goes into the prediction apparatus, and having a full structural theory is typically not viewed as essential. Longer-run forecasts, on the other hand, are often argued to be better if they rely on a more stripped-down and theory-oriented setting. We conjecture that epi-econ forecasting will have similar features. Over the very short run—say, dealing with one or two months ahead—the kinds of epi-econ IAMs advocated here are unlikely to be needed, or even optimal. However, for longer forecasts, we think they will be useful.

Second, one can systematically evaluate policy: perform counterfactual analysis. Thus, one uses any number of assumptions on the stance of policy, solves and simulates, and then compares. The outcome variables involve consumption and leisure, health, time use, etc., all by subgroup and over time. Here, one is primarily interested in how policy acts, even though these simulations—using the model as a lab for conducting experiments—of course also indicate the effectiveness, and hence desirability, of different options.

Finally, given that the epi-econ IAM is based on microeconomic theory, it can also be used for normative purposes. The normative use of the model works as follows. Step one is to realize that for any individual in the population, we have a utility-function representation allowing us to see which of two policy scenarios is best from this individual's perspective. Since the model involves the individual's consumption and leisure activity over time, including the expected life length, this evaluation naturally comes down to comparing present-value (expected) utility as of the beginning of time in the model. Step two is then to specify the (present-value, expected) utility outcomes for all individuals as a result of the policy experiment and then see whether one policy choice gives higher utility for all individuals; if it does, it is then a natural preferred policy. If there is no policy that is best for all, either one simply makes note of the distributional consequences (in utility space) of the policy. We can also—but do not need to—discuss what is "optimal" if we are willing to adopt a ***social welfare function***, detailing how the welfare of different people are traded off against each other. Typically such functions are designed to pay attention to inequality, but they do not have to.[15] It should be noted that the choice of a social welfare function is not based on economics but a philosophical, or political, one.

[15] In the case of climate change, a common parameter in a social welfare function is the "weight" on future generations: to what extent, if any, should the well-being of future generations be "discounted"?

6 Our epi-econ IAM

In this section we briefly outline our own simple epi-econ model. We began the construction of this model in early spring and it is currently complete; a paper will soon be made publicly available (Boppart et al. 2020). The model is a reaction to what we perceived as missing in early literature: we looked for a more explicit description of social activities (hence, a step toward sociology was needed), we thought time-use data would come in handy, and we found distinctions between more and less vulnerable people ("old" vs. "young") would be critical. We also felt that there was a need to build for the future, and hence construct a setting that could be applied also to different viruses than covid-19 (e.g., SARS or an ordinary flu). The model is, however, still quite stylized and should be viewed as a proof of concept.

Schematically, Figure 8 describes the people in the model and their activities.

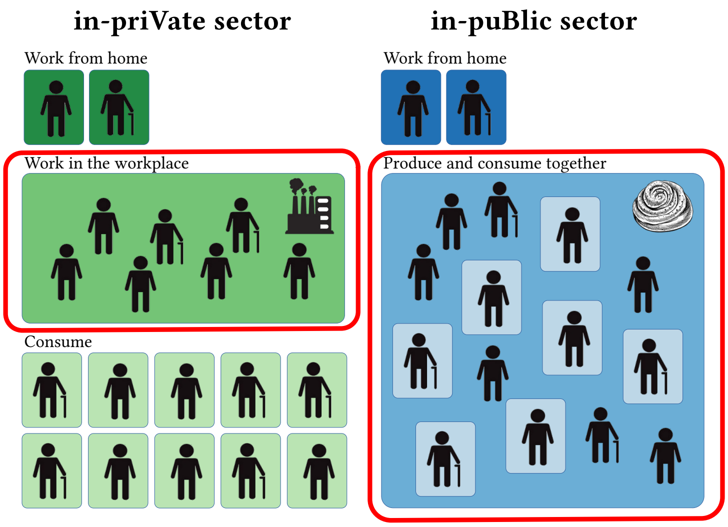

Figure 8: Illustration of the model. The areas marked in red are where the virus spreads.

There are two types of individuals, young and old. They can spend their leisure in the in-priVate sector (illustrated by the figures in the bottom left in the figure). When people spend their time on this type of leisure, e.g., watching Netflix, they do not interact with others and hence there is no risk of becoming infected. The goods and services used for the in-priVate leisure are produced in the workplace or at home (illustrated by the upper half of the left side of the figure, where people are, e.g., in the studio producing a Netflix show or in the Amazon warehouse shipping a new TV set). When people work in the workplace, they interact with their colleagues, and there is a risk of spreading or contracting the virus.

The right hand side of the figure illustrates the in-puBlic sector. In this sector, consumption and work take place jointly, and the virus can spread between those enjoying their leisure (e.g., customers in the restaurant) and those working in the sector (e.g., waiters in the restaurant). However, even in this sector there is a possibility (for at least some employees) to work from home without physically interacting with others.

In the beginning of the epidemic, a small fraction of the population is infected, while the majority are still susceptible, and no one has yet recovered. As the epidemic spreads, susceptible individuals become infected, while some infected individuals recover and a fraction of the infected dies. The infection fatality rate is higher for the old than for the young, and it increases if the hospitals become overcrowded. We assume full immunity, so that once an individual has recovered, he/she cannot be re-infected. Individuals in the model receive utility from leisure, from consumption, and from the intrinsic value of being alive.

We consider three versions of the model. One is a laissez-faire market equilibrium where people behave myopically: throughout time, they are fully unaware of the existence of the virus (or do not understand that their actions have any relation to the risk of getting infected). A second case is a market equilibrium where people instead are fully informed of the virus and rationally forward-looking. A third and final case is an "optimal allocation". The optimal allocation here is unique, as we adopt a family structure in our economic model: there is a large number of identical families and within each family a fraction of the family members are old, the remainder being young. (Also families are large, so that we can talk about fractions.) Thus, we adopt the concept of a *representative family* as the unique economic household in our model.[16] Hence, both in market allocations and that chosen by the planner, welfare is uniquely based on the representative family's (present-value, expected) utility outcome.

Because the social planner's problem is easier to express compactly, we base our presentation here on it.

6.1 Formal description of the social planner's problem

To formally state the social planner's problem, we first describe the state variables in the dynamic problem, and then state the maximization problem of the planner.

[16]This approach is taken in large parts of macroeconomics. Of course, it is not the only approach and there is now a large literature of "heterogeneous-agent" settings. These are more demanding to study but otherwise preferable.

State variables: S_t^i is the number of susceptible of type $i \in \{y, o\}$ (young/non-vulnerable and old/vulnerable, respectively) at time t; similarly I_t^i is the number of infected and R_t^i the number of recovered. The total population size in time t is $\sum_i \left(S_t^i + I_t^i + R_t^i\right)$, with initial conditions satisfying $I_0^y + I_0^o = 0.001, R_0^i = 0$ for $i \in \{y, o\}$ and $S_0^y + S_0^o = 0.999$. We assume that the fraction initially infected is equal in the young and the old population. In the absence of an epidemic, all young individuals live until period $\mathcal{T}^y$ and all old individuals live until period $\mathcal{T}^o < \mathcal{T}^y$.

Law of motion for the state variables: The law of motion for the state variables is then given by,

$$S_{t+1}^i = S_t^i - T_t^i \tag{5}$$

$$I_{t+1}^i = I_t^i(1 - \pi_r - \pi_{d,t}^i) + T_t^i \tag{6}$$

$$R_{t+1}^i = R_t^i + \pi_r I_t^i, \tag{7}$$

where

$$\pi_{d,t}^i = H(I_t^y + I_t^o). \tag{8}$$

The death rates $\pi_{d,t}^i$ is a function of the total number of infected individuals in the economy, since it depends on whether the hospital system is overcrowded or not.

The maximization problem: Individual utility is a function of consumption (c) and leisure (h). Leisure is of two sorts: leisure involving social interactions (e.g., restaurants or movies) and leisure not involving social interactions (watching TV, reading the newspaper); we use the indexes B (in-puBlic) and V (in-priVate) for these, respectively.

The consumption good is also of two sorts: goods and services consumed while being socially active (e.g., restaurants or movies), c_B, and goods and services consumed in-priVate (e.g., television), c_V.

A key assumption is that the planner cannot identify the individual health statuses but can compute the evolution of the total number of infected and recovered. Thus, the planner is not, for example, able to command only the recovered to specialize on the socially intense activities. Testing would provide this kind of information and by allowing testing we could evaluate the social benefits of tests.[17] Given these constraints on

[17]Theoretically, testing is somewhat cumbersome to analyze, since one would then ideally keep track not only of who has been tested but when. Thus, whereas the model we describe here amounts to an anonymous description of individuals, tests would create a large number of identifiable, and by health status

what can be accomplished, the planner chooses all the control variables and maximizes

$$\sum_{i \in \{y,o\}} \sum_{t=0}^{\mathcal{T}^i} \beta^t \left[(S_t^i + I_t^i + R_t^i) v(c_{B,t}^i, h_{B,t}^i, c_{V,t}^i, h_{V,t}^i) \right]$$

subject to the law of motion given by equations (5) to (8) and

$$1 = h_{B,t}^i + h_{V,t}^i + n_{Bh,t}^i + n_{Bw,t}^i + n_{Vh,t}^i + n_{Vw,t}^i \qquad i \in \{y,o\}, \tag{9}$$

$$\sum_{i \in \{y,o\}} \phi_t^i c_{j,t}^i = F_j(\phi_t^y n_{jh,t}^y, \phi_t^y n_{jw,t}^y, \phi_t^o n_{jh,t}^o, \phi_t^o n_{jw,t}^o), \qquad j \in \{B,V\}, \tag{10}$$

$$T_t^i = G(h_{B,t}^y, h_{B,t}^o, n_{Bw,t}^y, n_{Bw,t}^o, n_{Vw,t}^y, n_{Vw,t}^o, \Omega_t) \qquad i \in \{y,o\}, \tag{11}$$

$$\phi_t^i = S_t^i + I_t^i + R_t^i \qquad i \in \{y,o\} \tag{12}$$

for all t and non-negativity constraints for time and consumption quantities. Ω_t represents a complete description of the SIR state at time t, $\Omega_t = (S_t^y, I_t^y, R_t^y, S_t^o, I_t^o, R_t^o)$, i.e., the number of susceptible, infected, and recovered of young and old, respectively. Ω_0 is given exogenously.

Constraints (9) are the time constraints. Both young and old have one unit of time at their disposal, and can spend time on leisure (h_B and h_V) or working (n). Hours worked can take place in two different locations: either at home (indexed h) or in the workplace (indexed w) and each of these types of work can be carried out in two sectors: producing goods and services for the in-puBlic sector (e.g., restaurants or movies), or producing goods and services for the in-priVate sector (everything else).

Constraints (10) are the resource constraints. Production in both sector B and in sector V are a function of the hours worked from home and the hours worked in the workplace, for young as well as old.

Equation (11) tells us that the number of transmissions (i.e., newly infected) in each period depends on how much time is spent in the activities where the virus can spread, i.e., leisure of the in-puBlic type, and work in the workplace in either sector, as well as the complete SIR state of the economy. Finally, (12) states that the population in each period is the sum of susceptible, infected, and recovered individuals.

The solution to this maximization problem is the answer to the following question: if a social planner could decide on the actions of all individuals in the economy, how would it then allocate time and consumption in order to maximize the discounted sum of total utility, including both utility from consumption and leisure as well as the

fundamentally different groups.

intrinsic value of life?

Before the model can be used we need to parametrize and calibrate it.

6.2 Parametrization and calibration

To parametrize and calibrate the model includes specifying functional forms for $v(\bullet)$, $F(\bullet)$, $G(\bullet)$ and $H(\bullet)$. The calibration of the model is done in two separate steps. First, the economic parameters are calibrated to match a stationary environment with a population size of 1 (before an epidemic hit and people die). The population is divided into young (defined as age 15-60) and old (above 60). Second, for the epidemiological parameters we use data from the fast-growing literature on the covid-19 virus on the aggregation level needed for our model.

6.2.1 Parametrization and calibration of the economic side of the model

In this section, we highlight some important aspects of the calibration of the economic model. For a full description, we refer to Boppart et al. (2020).

Utility function We define $v(\bullet)$ as

$$v(c_B, h_B, c_V, h_V) = \underline{u} + u(c_B, h_B, c_V, h_V)$$

with $u(\bullet)$ being a nested constant-elasticity-of-substitution function of leisure and consumption of the two different types. We set $\underline{u}$ to be consistent with estimates of the value of a statistical life from the literature. We use two values for the value of a statistical life, one higher and one lower. For the higher value, we follow for instance The Environmental Protection Agency and the Department of Transportation in the US who estimate the value to \$11.5 million (see Greenstone and Nigam (2020)), which translates into a value of one *year* of life being roughly 11.4 times *yearly* per capita consumption in the US (Glover et al. 2020). This is a high value, relative to VSL numbers used in other contexts. We therefore also use a lower number, and set this to 4.0: a year of life is worth roughly four times annual consumption, a number that is well within the range of values discussed in Viscusi and Aldy (2003).

Production function Production in both the B and the V sector is given by a Cobb-Douglas function of capital and aggregate hours worked at home and in the workplace

$$F_j(n^y_{jh}, n^y_{jw}, n^o_{jh}, n^o_{jw}) \quad = \quad k_j^{\alpha} \tilde{n}_{jh}^{\nu} \tilde{n}_{jw}^{1-\alpha-\nu} \tag{13}$$

where, given our focus on short-run analysis, the sector-specific capital stocks k_B and k_V are fixed and total hours worked at home ($\tilde{n}_{jh}$) or in the workplace ($\tilde{n}_{jw}$) are aggregates of hours worked in each place by the young and by the old.

The parameters of the production functions are externally calibrated. We assume $\alpha = 1/3$, for standard reasons. In the data, working hours are on average distributed between work from home (n_h) vs. in the office (n_2) such that $n_w/n_h = 2.3$.[18] In the absence of an epidemic, the marginal products of n_h and n_w are equal, delivering $(1-\alpha-\nu)n_h = \nu n_w$. Hence we obtain that $\nu \approx 0.202$. The output elasticity of work that can be done from home is 0.202 whereas it is 0.465 for work that can only be done at the workplace.

We can use these values to assess how much production would be lost if n_w/n_h were forced to fall from 2.3 to, say, 1. Then output would be

$$\frac{3.3^{0.465} 3.3^{0.202}}{4.6^{0.465} 2^{0.202}} \approx 0.95,$$

i.e., output would fall by 5 percent. This is sizable, though not a huge amount. If n_w/n_h falls to 1/3 (2 hours worked at the workplace plus 6 hours worked from home out of a 8 hours workday) the output loss is 25 percent. We find these losses reasonable in magnitude, supporting a choice of $\nu = 0.202$.

Choice of remaining parameters In the choice of parameters we are restricted to actual time use data, as described in section 4.1. We want our model, in the absence of an epidemic, to replicate the time use patterns that we see in the economy. In particular, we need our model to replicate how much of the non-sleeping time is spent on leisure vs. work for young and for old respectively and how large fraction of our leisure is spent on the in-puBlic type and on the in-priVate type. We also need to capture the total amount of work done in the in-puBlic sector vs. in the in-priVate sector. Moreover, we want our model to display labor supply elasticities in line with the literature. Given these restriction, we can calibrate the parameters of the functions described above. For a full list of parameter choices we refer to Boppart et al. (2020).

[18]Based on calculations using ATUS, see Section 4.1.3.

6.2.2 Calibration of the epidemic

We now turn to the calibration of the epidemiological parameters for the SIR part of our model. First, we describe the functional form of $G(\bullet)$.

For each type $i \in \{y, o\}$, the number of new infections/transmissions is given by

$$T^i = T^i_B + T^i_V,$$

in other words, it is the sum of how many became infected in the B sector and how many became infected in the V sector.

The number of individuals of type i who got infected in the B sector is given by

$$T^i_B = \pi_B \frac{S^i\left(h^i_B + n^i_{B,w}\right)\sum_{j\in\{B,V\}} I^j(h^j_B + n^j_{B,w})}{\sum_{j\in\{B,V\}}\left[(S^j + I^j + R^j)(h^j_B + n^j_{B,w})\right]}.$$

Thus, it is proportional to *the time spent* in sector B either in the form of leisure or work in the workplace by susceptible individuals of type i, $S^i\left(h^i_B + n^i_{B,w}\right)$.[19] The number of transmissions also depends on the share of individuals, weighted by time spent in sector B, who are infected, $\frac{\sum_{j\in\{B,V\}} I^j(h^j_B+n^j_{B,w})}{\sum_{j\in\{B,V\}}\left[(S^j+I^j+R^j)(h^j_B+n^j_{B,w})\right]}$.

The number of infected in the V sector is analogously given by

$$T^i_V = \pi_V \frac{S^i n^i_{V,w}\sum_{j\in\{B,V\}} I^j n^j_{V,w}}{\sum_{j\in\{B,V\}}\left[(S^j + I^j + R^j)n^j_{V,w}\right]}.$$

with the difference that in the V sector, infections only take place among those working in the workplace.

The myopic model implies an R_0 that we restrict our key spread parameters π_B and π_V to match. To distinguish π_B from π_V amounts to drawing distinctions between the social interactions in the leisure activity (including for those who work in that activity) and the social interactions in the workplace in the production of the good used in the in-priVate composite. We consider all interactions equally contagious and set $\pi_B = \pi_V$.

The estimates of R_0 for covid-19 are uncertain and range between at least 1.4 and 3.9; we use 2.0 in our benchmark simulations. We simulate the simplest possible SIR model with a homogeneous population given this estimate of R_0 (and a recovery rate π_r to be specified below). This gives us a measure of the final number of recovered (which

[19]This is in contrast to the standard SIR model as described in section 3.1, where the transmissions only depend on number of individuals in each state, and not the time spent in contagious activities vs. non-contagious activities.

is 78%) if the epidemic were to play out unhindered. Thereafter we use the steady-state time allocations in our economic model and find the $\pi_B = \pi_V$ which give the same final number of recovered in the economy, were there no endogenous behavioral responses.[20] [21]

In line with Atkeson (2020) and Eichenbaum et al. (2020a) we set the average time from infection to recovery to be 18 days. This time also corresponds to the time from symptom to recovery in Glover et al. (2020). Since our model is daily, π_r, the recovery rate, is set to 1/18.

We assume that the death rate of the illness is an increasing function of the number of infected. $\pi^i_{d,t}$ is thus a logistic function for which the midpoint of the logistic curve–where the increase in death probability is the highest–occurs at the point where the hospitals are getting over-crowded. This point is assumed to happen when the fraction of infected in the population reaches $\hat{I}$. The current death rate in any time period is thus a function of the current total number of infected:

$$\pi^i_{d,t} = H(I_t) = \pi^i_{d,low} + \frac{\pi^i_{d,high} - \pi^i_{d,low}}{1 + e^{-k(I_t - \hat{I})}}$$

with I_t denoting the sum of the young and the old infected. Based on US data, there were 29.4 intensive care units (ICUs) per 100,000 people at the onset of the covid crisis so we assume one ICU per 3,400 people.[22] Further, we assume that three percent of the infected individuals require hospitalization, and, based on estimates for Sweden, that 29% of the hospitalized are in need of intensive care.[23] Taken together, this gives us an $\hat{I} = 1/(0.03 \times 0.29 \times 3400) \approx 0.034$. In other words, we assume that the death rate will quickly increase when the number of infected reaches 3.4% of the population.

The probability of dying (on a given day) conditional on being infected, when there is no over-crowding in the hospitals, is set to $0.001 \times 1/18$ for the young and $0.025 \times 1/18$ for the old, following Glover et al (2020). This means that the average infection fatality rate in the population—the share of infected who die—is 0.7%. When the health care system is completely overburdened, the probabilities are assumed to be twice as high

[20] Recall from the discussion of the simple SIR model above that the epidemic dies out before all individuals have become infected.

[21] The resulting epidemiological spread is very close but not exactly the same as the SIR model with homogeneous population simulated initially, since young and old have slightly different time allocations and slightly different death rates.

[22] Based on information from Society of Critical Care Medicine, downloaded from https://sccm.org/getattachment/Blog/March-2020/United-States-Resource-Availability-for-COVID-19/United-States-Resource-Availability-for-COVID-19.pdf?lang=en-US at June 24.

[23] Glover et al (2020) assume a hospitalization rate of 2% for the young and 12.5% for the old, which with our population shares would give a weighted average of 4.9%.

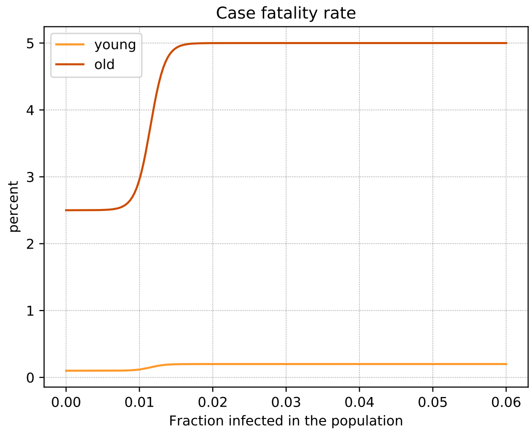

Figure 9: Infection fatality rate for young and old.

and thus set to $0.002 \times 1/18$ for the young and $0.05 \times 1/18$ for the old. The steepness of the curve, k, is set to 1000. The resulting mortality rates are shown in Figure 9.

The initial population is, as mentioned before, normalized to one. The number of people that are assumed to be initially infected is set to 0.001, evenly spread out among young and old.

A summary of the epidemiological parameters is given in Table 2.

6.3 Putting the model to use

In this section we contrast three different scenarios:

1. *A myopic scenario:* households do not realize that their time allocations affect their death probabilities (note that they might very well be aware of the epidemic, and realize that it is dangerous, without understanding the connection between their own actions and their own risk of becoming infected).

2. *A rational-expectations scenario:* all households have full information and do understand how their own actions affect their risk of getting infected. In particular, they understand that by choosing time use with less social interaction they reduce the risk that they become infected and die. To make a comparison with the planner solution, we assume that the individuals cannot observe whether they are susceptible, infected, or recovered. Although this may seem a bit odd, we argue that without perfect testing most of the infected would have a hard time knowing whether they have covid-19 or some other, less dangerous, health problem. In ad-

Parameter	Description	Value
Epidemic variables		
R_0	Spread factor standard SIR model	2.0
$\pi_B = \pi_V$	Spread factor economic model	0.24
π_r	Recovery rate	1/18
$\pi_{d,low}$	Death rate (before overcrowding)	$[0.001, 0.025] \cdot 1/18$
$\pi_{d,high}$	Death rate (when overcrowded)	$[0.002, 0.050] \cdot 1/18$
Health care system		
ι_h	Fraction of infected in need of hospitalization	0.03
ι_i	Fraction of hospitalized in need of ICU	0.29
ι_b	Inhabitants per ICU bed	3,400
x_0	Midpoint logistic function (fraction infected)	$1 / (\iota_h \cdot \iota_i \cdot \iota_b)$
k	Steepness parameter	1000

Table 2: Summary of epidemiological parameters. See text for description of sources and methodology.

dition, it is possible that they are infectious during part of the incubation period. In any case, perfect information about health status is not a better assumption unless a very good test becomes available.

3. *A social planner scenario:* where a social planner can decide on how to allocate time and consumption in order to maximize total utility

Below follow a number of key take-aways from the model.

The evolution of the epidemic is significantly different in the three scenarios In Figure 10, we show the evolution of the epidemic under the myopic market allocation, the rational-expectations market allocation, and the social planner's allocation under the assumption that no vaccine arrives. The epidemic under the myopic market allocation is close to standard SIR dynamics. The health system is overloaded, many young and old get infected, and the epidemic is essentially over because of herd immunity after 300 days.

The evolution of the epidemic under the rational-expectations market allocation may at first pass seem similar. The health system is overloaded and the epidemic is essentially over after 300 days. However, under rational expectations, few old become infected. The epidemic is primarily a risk for the old and under rational expectations they shift their behavior away from activities associated with infection risk. The young also do so, but to a much lesser extent both because their risk is lower and because the

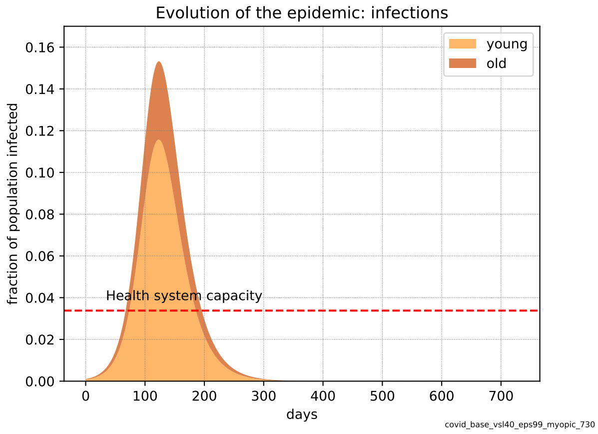

(a) The evolution of the epidemic under the myopic market allocation.

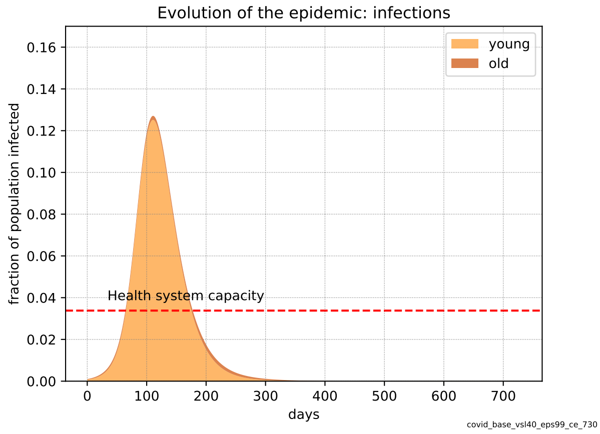

(b) The evolution of the epidemic under the rational expectations market allocation.

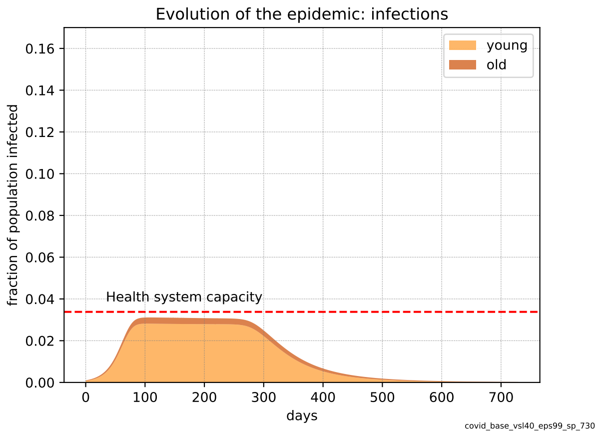

(c) The evolution of the epidemic under the social planner's allocation.

Figure 10: The evolution of the epidemic under the social planner's allocation and the two different market allocations.

labor-market wages give a compensating differential to the young.

Under the social planner's allocation, the evolution of the epidemic is qualitatively different. The social planner internalizes the effects of an overloaded health system and keeps infections below the threshold for overloading. Therefore, it takes a longer time to reach herd immunity, the epidemic is essentially over after 400 days.

In Figure 11, we show the evolution of the susceptible, infected, and recovered for the three scenarios. As expected, the final number of recovered is the lowest in the social planner scenario: a social planner ensures that the herd immunity threshold is reached with the smallest total number of people having become infected. In other words, a social planner avoids "over-shooting" in the number of infected and subsequently recovered.

The myopic scenario is very bad In the myopic scenario, where no adjustment of economic and social activities occurs, the epidemic has a very large cost in terms of human life, with approximately one percent of the population dying in the epidemic. The rational-expectations scenario significantly improves on this outcome, with only 0.21 percent of the population dying. Why? In the rational-expectations scenario, the individuals in the economy recognize the risk that they individually are facing (but ignore the indirect effect that their behavior has on others). The informed self-interest of the old is enough for them to drastically change their behavior, e.g., working from home and reduce their leisure in public to a bare minimum.

Optimal policy can be qualitative very different across cases In Subfigure 10c, we see that the optimal policy is to "protect the health care system", i.e., to keep the infection rate at a manageable level so that the hospital system is not overwhelmed. This conclusion depends crucially on the assumption that no vaccine or cure arrives. If we assume that a cure arrives after 6 months, the optimal policy is instead "full suppression", to keep infections very low while waiting for the vaccine to arrive. It is a quantitative question, depending on the potential arrival of a vaccine, but also on the other parameter values (e.g., the value of a statistical life) whether optimal policy is better described as "protect the health care system" or "full suppression".

A full assessment of policies requires valuing social leisure It turns out that, with our particular calibration configuration, the number of people who dies in the social-planner allocation and the rational-expectations allocation, is very similar: 0.21 percent of the population in the rational expectation scenario, and 0.20 percent

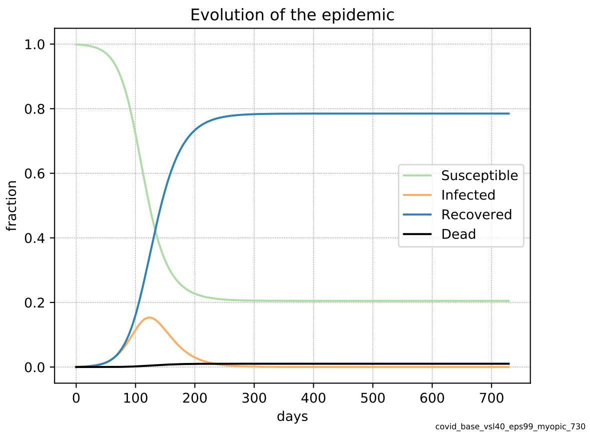

(a) SIR dynamics under the myopic market allocation.

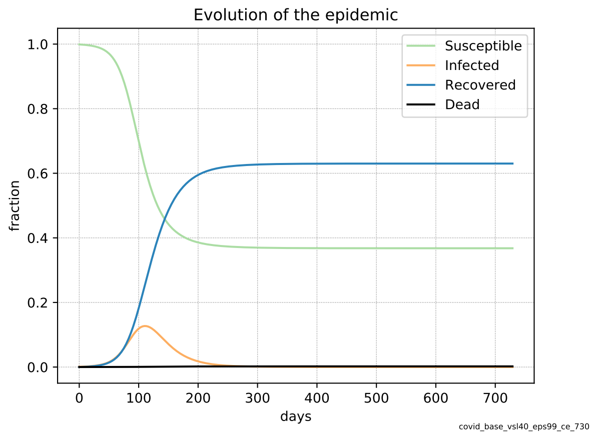

(b) SIR dynamics under the rational expectations market allocation.

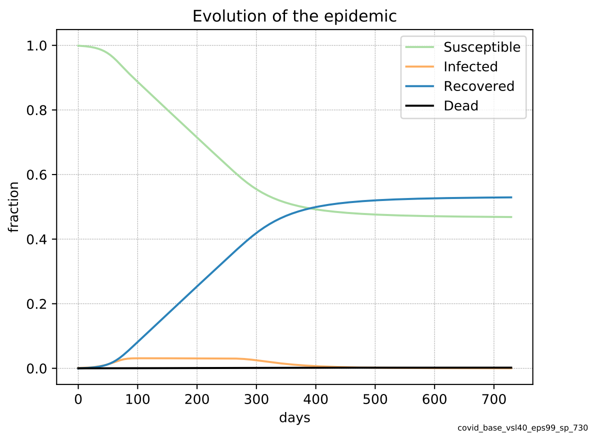

(c) SIR dynamics under the social planner's allocation.

Figure 11: SIR dynamics in the social planner's allocation and the two different market allocations.

in the social planner scenario.[24] The social-planner allocation implements a fall in production by almost 20 percent during the peak of the epidemic, and during the first year, output falls by 9.7 percent. The rational-expectations allocation only implements a fall in in output of five percent during the peak, and the fall during the first year is a mere 1.4 percent. How can the optimal policy generate both a larger fall in output and approximately the same number of deaths as the rational-expectations allocation and still be optimal? One answer is that the number of lost *years* of life is lower in the social planner scenario, due to the composition of deaths. However, another reason is as important. With optimal policy, the epidemic is kept in check, and as a result the old can spend some time on public leisure. By contrast, in the rational-expectations allocation, the infection risk is so high so that the old barely go outside at all. As a result of the improved leisure outcome—which is not captured in GDP measures—the old are *much* happier under the optimal policy.

A framework that frames epidemic policy as a tradeoff between economic indicators, such as output and unemployment, and lives saved risks missing the value of quality leisure. In our model, social leisure is explicitly valued by individuals and the cost of social isolation is incorporated into the analysis. This point thus speaks directly to the point raised in the introduction: even if the output paths of two economies looks similar, one can yield much higher welfare if it involves more non-work social interaction.

Intergenerational conflict of interest Since the health risks of an infection are different for young and old individuals, there is an intergenerational conflict of interest between generations. In the rational-expectations scenario, the young are living their lives almost as normal while the old are reducing their social leisure to an absolute bare minimum. As a result, the young suffer little during the epidemic while the old are sacrificing a lot. By contrast, in the social-planner allocation, the young also reduce their social activities in order to reduce the spread of the epidemic. This primarily helps the old. There is therefore an intergenerational conflict of interest. Optimal policy reduces the period utility of the young to the benefit of the old, which also leads

[24] Although almost the same number of individuals die in the two scenarios, the composition of deaths is very different. In the rational expectation scenario 54% of the deaths are in the young population. In the social planner scenario, 24% of the deaths are in the young population. As we saw in Figure 11, the final number of recovered is lower in the social planner scenario than in the rational expectation scenario. However, even though the total number of individuals who have once been infected is lower in the social planner scenario, total number of deaths is equal since there were more old individuals among the infected, and the old have a higher death rate. Thus, a social planner saves young lives compared to a rational expectations scenario, which translates into more years of life saved.

to slightly more deaths among the old, and slightly less deaths among the young.

Cross-virus validation of the framework It is crucial to test the model and its assumption. One way to do so is to test the framework against other viruses, and we do this by simulating a seasonal flu. This exercise helps validating our assumptions, especially the critical assumption of the intrinsic value of a statistical life. There are numerous estimates in the literature, and we do not (and should not!) take a stance on the correct value. However, in this model a value of a statistical life from the higher end of the spectrum implies that a social planner would want to slow the economy substantially to stop the spread of a seasonal flu: it would want to lower output by 5% during a critical quarter. This, as far as we can tell, is not how actual policy makers have reacted historically. Thus, although we cannot say whether a chosen value of a statistical life is the correct one, we can say that a value from the lower range is more in line with observed policy actions. When we test our model by simulating SARS, a disease both more contagious and more lethal than covid-19, we get that a social planner would impose a very strict lock-down even with the lower value of life (assuming a year until we can end the epidemic exogenously).

7 Concluding comments

In this report, we have focused on describing the interaction between the economy and an epidemic. We have disregarded various market imperfections and frictions that can motivate massive economic interventions such as the ones we have seen during the pandemic. Our neglect here is not because we think such policies are unimportant. On the contrary, these policy measures have likely prevented the covid-19 crisis from turning into a full-blown depression. However, the analysis of such policies can be evaluated using standard economic models by considering the changes in policy and individual behavior as an exogenous shock to the economy. In such an analysis, economic policy plays the second fiddle, responsible for minimizing the negative economic side effects of the shocks to the economy induced by the pandemic.

The pandemic shocks are best modeled as temporary supply-side restrictions. People spend less time in restaurants or at sports arenas not because their preferences have changed so they demand less of these services but instead because during the pandemic, these establishments cannot provide the services in a safe way.

Our report sketches a framework where economics plays another role than the second fiddle: though it does not play first fiddle, it plays as an equal partner with epidemiology. The aim of this partnership is to find policies that are optimal from a joint perspective. The framework is fully operational but not yet ready for full quantitative analysis. However, we do think it already brings us some important insights.

First, epidemiologists and social scientists must talk to each other about policy. Epidemiologists should be asked to describe how different policies—recommendations as well as regulations—are expected to affect contagion. Which type of behavioral changes are targeted and in which groups? How is this supposed to change the transmission of the infection? These questions should be answered quantitatively, but also a rough classification of measures as having large or small effects or a broad vs. a narrow focus is valuable. Given such information, economists and other social scientists can evaluate the consequences of the different policies in other dimensions than the epidemiological one. Furthermore, social scientists are trained to analyze how a targeted change in behavior, in the marketplace but also outside it, can be implemented with efficient policy tools.

Second, the set of policies to be discussed and evaluated must be broad enough. Finding a good trade-off between cost and benefits of policies against the pandemic requires that many different policies are evaluated. It is not sufficient for policymakers to only consider policy proposals based on one perspective, be it epidemiological or economic. Instead, experts with different perspectives must be asked to help provide counter-factual analyses. What would happen if something else than the proposed policy were to be implemented? Due to the large uncertainty, different scenarios, i.e., predictions under different assumptions, need to be presented.

Third, the economic perspective must also be wide. There is a risk that the economic consequences of policy is studied with a focus on what is measured only by what is given in our national accounts. But many epidemic policies are likely to have effects that are not visible in GDP or the government budget. In our framework, we emphasize the value of leisure which, which in ordinary language for example can be thought of in terms of seeing friends or having family over for Christmas, attending cultural or sports events, or merely traveling. The cost of distorting the use of leisure is in principle not less important than the cost of distorting time spent working for a wage. Education is another use of time where output is not measured in GDP. However, we have good reasons to believe that the educational system produces human capital with an economic value per student equal in magnitude to what is produced per worker in the marketplace. Thus, policies that negatively affect the quality of education can be very costly despite not affecting the current flow of GDP.

Given all the work that is taking place right now, we are confident that powerful epi-econ IAMs will be available soon and that these tools will prove very valuable for actual policy-making. A final point here is that a strength of explicit frameworks is that they can be communicated. Clear communication of the basis for decisions is important, for it allows scrutiny and constructive critique. If policy decisions are made without reference to systematic frameworks of evaluation there is a risk that the credibility in government and government agencies will slowly wane.

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2020). A multi-risk sir model with optimally targeted lockdown. Technical report, National Bureau of Economic Research.

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. E. Ozdaglar (2020). Testing, voluntary social distancing and the spread of an infection. Technical report, National Bureau of Economic Research.

Adda, J. (2007). Behavior towards health risks: An empirical study using the "Mad Cow" crisis as an experiment. *Journal of Risk and Uncertainty 35*(3), 285–305.

Alon, U., T. Baron, Y. Bar-On, O. Cornfeld, R. Milo, and E. Yashiv (2020). COVID-19: Looking for the exit. Technical report, working paper.

Alvarez, F. E., D. Argente, and F. Lippi (2020). A simple planning problem for covid-19 lockdown. Technical report, National Bureau of Economic Research.

Atkeson, A. (2020). What will be the economic impact of covid-19 in the US? Rough estimates of disease scenarios. Technical report, National Bureau of Economic Research.

Aum, S., S. Y. T. Lee, and Y. Shin (2020). Inequality of fear and self-quarantine: Is there a trade-off between GDP and public health? Technical report, National Bureau of Economic Research.

Bethune, Z. A. and A. Korinek (2020). Covid-19 infection externalities: Trading off lives vs. livelihoods. Technical report, National Bureau of Economic Research.

Bodenstein, M., G. Corsetti, and L. Guerrieri (2020). Social distancing and supply disruptions in a pandemic.

Bognanni, M., D. Hanley, D. Kolliner, and K. Mitman (2020). Economic activity and covid-19 transmission: Evidence from an estimated economic-epidemiological model.

Boppart, T., K. Harmenberg, J. Hassler, P. Krusell, and J. Olsson (2020). An epi-econ IAM.

Brotherhood, L., P. Kircher, C. Santos, and M. Tertilt (2020). An economic model of the covid-19 epidemic: The importance of testing and age-specific policies.

Chan, T. Y., B. H. Hamilton, and N. W. Papageorge (2016). Health, risky behaviour and the value of medical innovation for infectious disease. *The Review of Economic Studies 83*(4), 1465–1510.

Chang, R. and A. Velasco (2020). Economic policy incentives to preserve lives and livelihoods. Technical report, National Bureau of Economic Research.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020a). The macroeconomics of epidemics. Technical report, National Bureau of Economic Research.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020b). The macroeconomics of testing and quarantining. Technical report.

Farboodi, M., G. Jarosch, and R. Shimer (2020). Internal and external effects of social distancing in a pandemic. Technical report, National Bureau of Economic Research.

Garibaldi, P., E. R. Moen, and C. A. Pissarides (2020, April). Modelling contacts and transitions in the sir epidemics model. *Covid Economics 5*, 1–20.

Geoffard, P.-Y. and T. Philipson (1996). Rational epidemics and their public control. *International economic review*, 603–624.

Giagheddu, M. and A. Papetti (2020). The macroeconomics of age-varying epidemics.

Giannitsarou, C., S. Kissler, and F. Toxvaerd (2020). Waning immunity and the second wave: Some projections for sars-cov-2.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Rios-Rull (2020). Health versus wealth: On the distributional effects of controlling a pandemic. *PIER Working Paper* (20-014).

Greenstone, M. and V. Nigam (2020). Does social distancing matter? *Becker Friedman Institute for Economics Working Paper* (2020-26).

Greenwood, J., P. Kircher, C. Santos, and M. Tertilt (2019). An equilibrium model of the african HIV/AIDS epidemic. *Econometrica 87*(4), 1081–1113.

Hall, R. E., C. I. Jones, and P. J. Klenow (2020). Trading off consumption and covid-19 deaths. Technical report, National Bureau of Economic Research.

Jones, C. J., T. Philippon, and V. Venkateswaran (2020). Optimal mitigation policies in a pandemic: Social distancing and working from home. Technical report, National Bureau of Economic Research.

Kapicka, M. and P. Rupert (2020). Labor markets during pandemics. *Manuscript, UC Santa Barbara*.

Kaplan, G., B. Moll, and G. Violante (2020). The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US. Technical report, National Bureau of Economic Research.

Kermack, W. O. and A. G. McKendrick (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), 700–721.

Kremer, M. (1996). Integrating behavioral choice into epidemiological models of AIDS. *The Quarterly Journal of Economics 111*(2), 549–573.

Krueger, D., H. Uhlig, and T. Xie (2020). Macroeconomic dynamics and reallocation in an epidemic. Technical report, National Bureau of Economic Research.

Moser, C. A. and P. Yared (2020). Pandemic lockdown: The role of government commitment. Technical report, National Bureau of Economic Research.

Piguillem, F. and L. Shi (2020). Optimal covid-19 quarantine and testing policies.

van Vlokhoven, H. (2020). Policy during an epidemic with super-spreaders. *Available at SSRN 3600441*.

Viscusi, W. K. and J. E. Aldy (2003). The value of a statistical life: A critical review of market estimates throughout the world. *Journal of Risk and Uncertainty 27*, 5–76.