

WORKING PAPER 2101

Pandemics and Aggregate Demand: a Framework for Policy Analysis

Peter Flaschel, Roberto Veneziani, Daniele Tavani and Giorgos Galanis

January 2021



Pandemics and Aggregate Demand: a Framework for Policy Analysis

Peter Flaschel*, Giorgos Galanis†, Daniele Tavani‡, Roberto Veneziani§

January 12, 2021

Abstract

This paper studies the interaction between epidemiological dynamics and the dynamics of economic activity in a demand-driven model in the structuralist/post-Keynesian tradition. On the one hand, rising aggregate demand increases the contact rate and therefore the probability of exposure to a virus. On the other hand, rising infection lowers aggregate demand because of reduced household spending. The resulting framework is well-suited for policy analysis through numerical exercises. We show that, first, laissez-faire gives rise to sharp fluctuations in demand and infections before herd immunity is achieved. Second, absent any restrictions on economic activity, physical distancing measures have rather limited mitigating effects. Third, lockdowns are effective, especially at reducing death rates while buying time before a vaccine is available, at the cost of a slightly more pronounced downturn in economic activity compared with alternative policies. This casts some doubt on the so-called “lives versus livelihood” policy trade-off. However, we also highlight the importance of policies aimed at mitigating the effects of the epidemic on workers’ income.

JEL Classification Codes: I1 (Health); E6 (Macroeconomic Policy, Macroeconomic Aspects of Public Finance, and General Outlook); E25 (Aggregate Factor Income Distribution); E12 (Keynes, Keynesian, Post-Keynesian, Modern Monetary Theory); H0 (Public economics, General)

Keywords: pandemic, aggregate demand, distribution, public policy.

*Department of Economics, Bielefeld University, Bielefeld, Germany. Email: pflaschel@uni-bielefeld.de

†Institute of Management Studies, Goldsmiths University of London, New Cross, London SE14 6NW, U.K. Email: g.galanis@gold.ac.uk

‡Department of Economics, Colorado State University. 1771 Campus Delivery, Fort Collins, CO, United States. Email: daniele.tavani@colostate.edu

§School of Economics and Finance, Queen Mary University of London, Mile End Road, London E1 4NS, U.K. Email: r.veneziani@qmul.ac.uk

1 Introduction

The SARS-CoV-2 pandemic is without a doubt one of the main policy concerns at the global level. It is almost a platitude to note that it is, and should be, a central focus for the activity of governments around the globe, and most social and economic measures currently under discussion revolve around the public response to the pandemic. Implementing the right set of policies – *both* public health measures *and* economic stimulus and relief packages – is important in order to avoid major losses in social welfare. As is evident from international comparisons, an erratic, uncoordinated, lax response to the virus has immediate, major negative social and economic repercussions compared with approaches informed by science and sound economic theory.

But the proposal of effective and equitable policy measures to be adopted here and now in order to tackle the emergency is only *one* reason to reflect on the SARS-CoV-2 pandemic. For, all evidence suggests that the virus is here to stay: even with a vaccine, it will take time to build herd immunity, and social and economic life will likely not go back to normal for months to come. Perhaps more importantly, SARS-CoV-2 is not the first coronavirus to spread globally, and it will not be the last. Any lessons learned from it are going to be precious in the next decades.

It also seems clear that the pandemic will have long-run structural effects. It is too early to paint a complete picture but some signs are already clear: while certain sectors are severely hit, others have gained significantly, and at least some of these structural changes are likely to be irreversible due to behavioural changes, or fixed costs. Perhaps more importantly in an era of staggering inequality, the pandemic is likely to have major distributive implications via its effect on labour markets, both at the micro and at the macroeconomic level. For example, once certain emergency government measures (such as bans on layoffs or major furlough programmes) are lifted, the ensuing unemployment may hit wages, thus further skewing the functional distribution of income.

Finally, the SARS-CoV-2 pandemic represents a unique natural experiment in comparative political economy. Faced with a similar challenge, different socio-economic structures have responded very heterogeneously, leading to significant divergences in outcomes, which are likely to persist.

In this paper, we develop a simple, flexible framework to analyse the joint dynamics of epidemiological variables and economic magnitudes, and to appraise the various policies adopted in different countries. The model is made of two building blocs. We begin by adopting the SEIR (Susceptible-Exposed-Infectious-Recovered) framework initiated from the work of [Kermack and McKendrick \(1927\)](#), which is a workhorse in epidemiology (for an overview of this type of models see [Hethcote, 2000](#)). Along the lines of a fast growing literature on the topic (for example see [Goenka et al., 2014](#); [Eichenbaum et al., 2020](#); [Atkeson, 2020](#)), we extend this by explicitly linking economic activity with the infection rate. This extension constitutes the first building bloc of our approach.

The second building bloc is the main conceptual innovation of the paper: following the post-Keynesian and structuralist tradition and drawing on [Chiarella and Flaschel \(2000\)](#), [Chiarella et al. \(2005\)](#), and [Flaschel et al. \(2018\)](#), we model economic activity by focusing on the rate of utilisation of installed capital stock. Because the epidemic has a negative impact on the population, and therefore on the labour force, we postulate a negative feedback of the number of infected people on the

dynamics of economic activity, and thus of capital utilisation.

Together the two blocs capture the dynamic interaction of epidemiological variables and economic activity in the absence of government intervention. To be sure, the description of the relevant economic mechanisms is rather stylised as we are capturing a whole host of economic factors by focusing only on one variable, while eschewing an explicit behavioural analysis. In this respect, our approach differs from much of the literature in economic epidemiology where agents are assumed to have full information on the contagion dynamics and its effects on health (for example, see [Sethi, 1978](#); [Chen, 2012](#); [Chen and Toxvaerd, 2014](#); [Toxvaerd, 2019](#); [Atkeson et al., 2020](#); [Eichenbaum et al., 2020](#); [Toxvaerd, 2020](#); [Rowthorn and Toxvaerd, 2020](#)). Our framework also differs from behavioural epidemiological models which assume boundedly rational individuals with limited information but do not model explicitly economic activity ([Manfredi and D’Onofrio, 2013](#); [Eksin et al., 2019](#); [Di Guilmi et al., 2020](#); [Galanis et al., 2020](#)). Our model contributes to a growing literature which analyses the epidemiological/economic interactions while not assuming fully informed rational agents ([Basurto et al., 2020](#); [Bellomo et al., 2020](#); [Caiani et al., 2021](#); [Razmi, 2020](#)).¹

The reduced form description of the dynamic interaction between epidemiological factors and economic variables – which may be dubbed the SEIR- u model – provides a simple, flexible framework to tackle a range of policy questions, and to compare different approaches to the pandemic. Furthermore, the focus on capital utilisation allows us to provide a natural way to capture lockdown measures as hard ceilings on economic activity leading to a forced reduction in the use of productive endowments. Finally, this framework is well suited to study some of the possible effects of the pandemic on the functional distribution of income.

The model provides a stark but neat exposition of the main policy tradeoffs faced by governments aiming to control the epidemic without excessive disruptions of economic activity. First, a laissez-faire strategy results in sharp fluctuations both in aggregate demand and in infections before the herd immunity threshold is crossed. Second, physical distancing without hard restrictions on economic activity is unlikely to tame the dynamics of infection: we show that, under conservative parameterisations of the model, physical distancing measures have rather limited mitigating effects on the epidemic and thus on the epidemic-induced recession. Third, and similar to [Kaplan et al. \(2020\)](#), the model is naturally suited to model lockdowns as hard caps on economic activity as opposed to reductions in contacts. We show the effectiveness of lockdowns at reducing infection – and consequently the number of pandemic-related deaths – while “buying time” before mass vaccination or treatments are rolled out. These results are robust to several modifications of the model, including an extension that features the dynamics of the functional distribution of income. In particular, we show the importance of accompanying public health measures with robust countercyclical policies – such as employment protection, unemployment insurance, direct transfers, or payroll subsidies to firms, – in order to mitigate the adverse distributional consequences of the epidemic.

Regarding the organisation of the paper, after introducing the full model in section [2](#), we study

¹This does *not* mean that the model is inconsistent with reasonable assumptions on individual behaviour. As shown in Appendix [A](#), behavioural foundations can be provided to our main equations.

a stripped down version that reduces to two-dimensions – utilisation and infections – and is useful to think through the short-run evolution of the pandemic. This simplified version is similar to the model studied by Gallo and Barbieri-Goes (2020), who focused instead on the short-run interaction between the pandemic and the unemployment rate. We then turn to simulating alternative policy scenarios in the full epidemiological-economic model in section 3. Section 4 studies the inclusion of distributional dynamics as well as a variety of policy exercises. Section 5 summarises the battery of robustness checks that we have run on the model. Section 6 concludes.

2 The SEIR- u model: Epidemic dynamics and economic activity

We first describe the complete model featuring the dynamic interaction between economic activity – as measured by the rate of utilisation of installed capital stock – and infections. The framework has a simple predator-prey structure: as economic activity increases, the virus spreads, although with a delay. The rise in infections depresses economic activity through a reduction in aggregate demand. The resulting decline in economic activity, in turn, results in a decline in infections.

2.1 Epidemic dynamics

Our starting point is a baseline SEIR model. Consider a large but constant population of N agents who, at any point in time t , are partitioned into four *compartments* based on their health condition: the susceptible, the exposed (who are infected but not yet infectious), the infectious, and the recovered, whose cardinalities are, respectively, $S(t)$, $E(t)$, $I(t)$, and $R(t)$. Formally, at all t ,

$$N = S(t) + E(t) + I(t) + R(t).$$

The dynamics of the four compartments is given by the following differential equations. First, the number of susceptible individuals changes over time as they come into contact with the infectious:

$$\dot{S}(t) = -\frac{\beta(t)S(t)I(t)}{N}, \quad (1)$$

where $\beta(t)$ captures the average number of contacts per person per time, multiplied by the probability of disease transmission in a contact between a susceptible and an infectious subject.

Once infected, susceptible individuals become exposed. If an exposed individual remains exposed for a period $1/\sigma$ before becoming infectious, σ is the rate at which an exposed individual becomes infectious. Thus, the evolution of $E(t)$ is:

$$\dot{E}(t) = \frac{\beta(t)S(t)I(t)}{N} - \sigma E(t). \quad (2)$$

²A complete description of all robustness checks and the code of the simulations can be found in the online Addendum, available at www.danieletavani.com/code.

³Although we do not explicitly include supply considerations in the paper, supply bottlenecks are also responsible for reductions in economic activity. Our argument is not affected by such consideration.

If an individual remains infectious on average $1/\gamma$ periods, – and therefore recovers (or dies) at a rate γ per unit of time, – the evolution of the infectious compartment of the population satisfies:

$$\dot{I}(t) = \sigma E(t) - \gamma I(t). \quad (3)$$

Finally, the compartment of recovered individuals evolves as follows:

$$\dot{R}(t) = \gamma I(t). \quad (4)$$

Because N is constant, we normalise it to 100. Accordingly, we interpret $I(t)$ as the percentage of infectious individuals in the population – or *infection rate*, for short – at t .

A crucial element of our analysis is the observation that the average number of contacts at any point in time depends, among other factors, on the overall level of social and economic activity, and aggregate demand in particular. For this reason we express $\beta(t)$ as an increasing function of capacity utilisation $u(t) \equiv \frac{Y(t)}{K(t)}$, where $Y(t)$ is aggregate income (spending) and $K(t)$ is the total capital stock⁴ $\beta(t) = \psi(u(t))$. Given that $\beta(t)$ captures the probability of becoming exposed, we adopt a linear specification for ψ ⁵:

$$\beta(t) = \beta_0 u(t), \quad (5)$$

where $\beta_0 > 0$. The dynamics of the susceptible and the exposed then modify as follows:

$$\dot{S}(t) = -\frac{\beta_0 u(t) S(t) I(t)}{N}, \quad (6)$$

$$\dot{E}(t) = \frac{\beta_0 u(t) S(t) I(t)}{N} - \sigma E(t). \quad (7)$$

Let $r(t) \equiv \beta(t)/\gamma$ be the so-called *simple* reproduction rate. At the beginning of the epidemic (that is with $S \approx N$), the short run condition for the growth rate of infections to be negative requires that the simple reproduction rate be less than one, that is: $u(t) < \frac{\gamma}{\beta_0}$ ⁶. We assume that $\bar{u} > \gamma/\beta_0$, which reflects the idea that the number of contacts associated with a normal level of economic activity fosters the diffusion of the virus, absent mitigation policies. In other words, economic activity does make a difference for the evolution of the epidemic.

2.2 Economic dynamics

Even in the absence of physical distancing measures or harder lockdown policies imposed by the government, economic activity is affected by the number of infectious individuals owing both to

⁴Capital utilisation, defined as the income-capital ratio, is a common proxy for aggregate demand in demand-driven growth models in the structuralist and post-Keynesian literature, and we find it the most intuitive concept. Nonetheless, all of our conclusions continue to hold if $u(t)$ is defined as the ratio of actual output over potential output, or capital used productively over installed capacity.

⁵All of our main results continue to hold under a range of alternative assumptions on ψ (see section 5.2 below).

⁶We analyse the dynamics of the reproduction rate in the various scenarios in Appendix C.

the direct effect of the virus on the labour force, and broader supply and demand effects. For example, data from the [Bureau of Economic Analysis](#) show that in the US real personal consumption expenditure declined by 0.4 percent, while disposable personal income declined by 1.2 percent between October and November 2020, a period of rising infection without hard lockdowns. We capture the repercussions of the epidemic on economic activity by assuming that a higher number of infections tends to reduce the growth rate of national income, denoted by $\hat{Y}(t)$:⁷

$$\hat{Y}(t) = n - \alpha I(t), \quad (8)$$

where n can be interpreted as capturing the long-run (trend) growth rate of the economy absent the epidemic, assumed to be exogenous throughout.

Next, abstracting from depreciation, we model investment through a simple accelerator equation as a function of capital utilisation as well as featuring a long-run growth component:⁸

$$\hat{K}(t) = n + \delta(u(t) - \bar{u}). \quad (9)$$

Therefore, given equations (8)-(9), the dynamics of the utilisation rate is given by $\hat{u}(t) = \hat{Y}(t) - \hat{K}(t)$, from which we have:

$$\dot{u}(t) = [\delta(\bar{u} - u(t)) - \alpha I(t)] u(t). \quad (10)$$

Equations (3), (4), (6), (7), and (10) form a five-dimensional dynamical system that links up the evolution of the epidemic to the dynamics of economic activity. The epidemiological bloc of the model is widely understood, and the integration of economic dynamics is quite straightforward. The complete model can be fruitfully analysed numerically, which we will do in section 3 below. From an analytical standpoint, the turning point of the epidemic is the crossing of the herd immunity threshold, which in terms of the infectious population is given by the inverse of the simple reproduction rate at the onset of the disease: $r(0)^{-1}$. After herd immunity is reached, infections converge to zero, and economic activity returns gradually to normal.

From a policy standpoint, however, the fact that the epidemic will eventually disappear is of limited interest if in the meantime significant disruption in economic activity and a large number of deaths occur. Our focus will therefore be on the short-to-medium run. In order to illustrate the core economic-epidemiological interaction and get some intuitions about the short-to-medium run behaviour of the model, we present a useful special case that strips down the epidemic bloc into a single equation describing the evolution of the infectious compartment of the population, so that the model reduces to two dimensions. This special case has the advantage of being analytically tractable and of neatly illustrating the perils of a prolonged epidemic-induced recession before the

⁷A behavioural microeconomic foundation for equation (8) is provided in Appendix A.

⁸The assumption that in the long-run capital and GDP grow at the same rate ensures that the economy is in balanced growth, at a rate equal to n , at a steady state with normal utilisation and no infections. This simplification rules out so-called Harrodian instability in the labour market, with employment either exploding or going to zero, and it allows us to focus on the short-run fluctuations induced by the epidemic.

population reaches the threshold required for herd immunity.

2.3 A short-run model of infection and aggregate demand

The simplest version of our framework can be obtained under two simplifying assumptions: (a) removing the difference between exposed and infectious, and (b) assuming that the number of susceptible individuals equals the total population. These two assumptions give rise to a stripped down two-dimensional dynamical system given by equation (10) above and

$$\dot{I}(t) = [\beta_0 u(t) - \gamma]I(t), \quad (11)$$

that can be useful when thinking about the short-run implications of the pandemic before herd immunity is achieved.⁹ The dynamical system has two non-trivial steady states:¹⁰

(i) A *no-infection* steady state with utilisation at its normal rate:

$$u^{NI} = \bar{u}; \quad I^{NI} = 0; \quad (12)$$

(ii) An *endemic* steady state with a positive number of infectious agents and below-normal utilisation.¹¹

$$u^E = \frac{\gamma}{\beta_0}; \quad I^E = \frac{\delta(\beta_0 \bar{u} - \gamma)}{\alpha \beta_0}. \quad (13)$$

2.3.1 Local Stability Analysis

The Jacobian matrix of the dynamical system evaluated at a steady state (I^{ss}, u^{ss}) is:

$$J(I^{ss}, u^{ss}) = \begin{bmatrix} \beta_0 u^{ss} - \gamma & \beta_0 I^{ss} \\ -\alpha u^{ss} & -\delta u^{ss} \end{bmatrix}.$$

At the no-infection steady state $(I^{NI} = 0, u^{NI} = \bar{u})$, our assumptions imply $\beta_0 \bar{u} - \gamma > 0$ and therefore the determinant is negative: the two eigenvalues of the matrix are of opposite sign, which implies that the corresponding steady state is a saddle point.¹²

At the endemic steady state, $\beta_0 \bar{u} - \gamma = 0$ and therefore the matrix has negative trace and positive determinant, indicating local stability instead. The eigenvalues of the matrix, denoted by $(\varepsilon_1, \varepsilon_2)$, are the solution to the characteristic equation:

$$\varepsilon_{1,2} = -\frac{\delta u^E}{2} \pm \frac{\sqrt{(\delta u^E)^2 - 4\beta_0 I^E \alpha u^E}}{2}. \quad (14)$$

⁹This special case has similar properties to the model recently studied in Gallo and Barbieri-Goes (2020). The difference lays in our focus on utilisation (instead of unemployment) as the main measure for economic activity.

¹⁰There also exists a trivial steady state with $u^T = I^T = 0$, which can be safely ignored.

¹¹Both properties follow from the assumption that $\gamma/\beta_0 \in (0, \bar{u})$.

¹²If β_0 was sufficiently low, or γ was sufficiently high, then $\beta_0 \bar{u} - \gamma < 0$, the determinant would be positive and the trace would be negative, ensuring local stability. However, this scenario implies that the virus has an inherent tendency to disappear under a business as usual scenario, which is most definitely not the case for SARS-CoV-2.

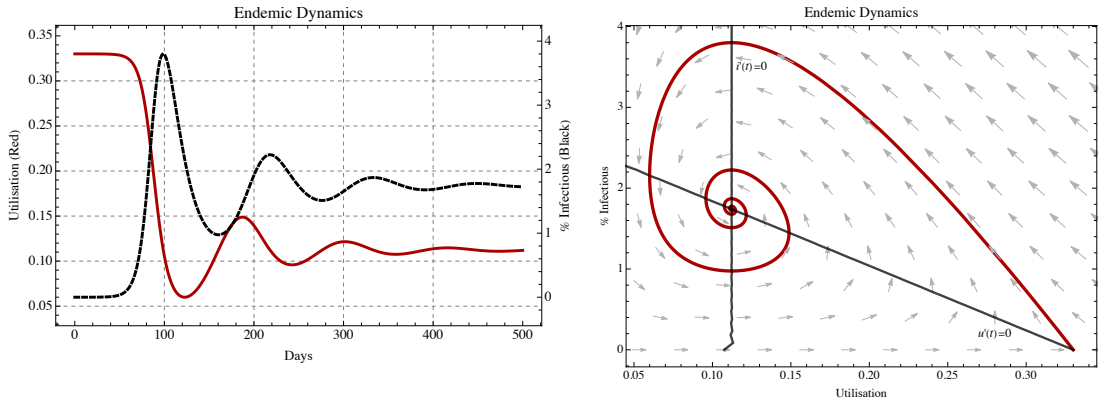
Therefore, at the endemic steady state, cyclical convergence occurs provided the discriminant in equation (14) be negative, so that the eigenvalues have imaginary parts. Simple manipulation reveals that a necessary and sufficient condition is

$$\frac{u^E}{\bar{u}} < \frac{1}{1 + \frac{\delta}{4\beta_0}}. \quad (15)$$

Under rather general assumptions on the main parameters (including the values used in our numerical implementation, see Table 1 below), the latter condition holds: convergence to the endemic steady state is therefore cyclical. The economic intuition for the result has to do with the fact that a rising aggregate demand increases the economy-wide contact rate, which in turn leads to rising infection. But the higher infection rate depresses aggregate demand, which lowers the contact rate and reduces infections. Fluctuations in aggregate demand and infections are ultimately damped, however, and the dynamics converges to the endemic steady state.

Figure 1 illustrates the dynamics of the two main variables for a representative set of parameters:¹³ the left panel plots the two variables against time. The right panel presents the phase plot of the dynamics of the economy.

Figure 1: Baseline 2D simulation: laissez-faire scenario.



3 SEIR- u model: Simulations and policy exercises

The complete epidemiological-economic model comprises equations (3), (4), (6), (7), and (10). Given the high dimensionality of the system it is difficult to gain insights on the interaction between epidemiological variables and economic activity analytically, especially in the study of out of steady state dynamics and policy exercises. We therefore turn to numerical methods.

Because our aim is to examine different approaches to the pandemic, and their economic implications, we do not calibrate the model to the data of any specific country. Rather, we choose a set of parameters that reflect what may be considered to be representative, or average values across a

¹³See Table 1. Recall that in the 2D model, $I(t) = E(t)$ all t .

range of countries based on the available evidence reported in the literature¹⁴ We then run a large number of checks in order to verify that our main conclusions are robust to significant perturbations of the parameters (see section 5.1 below).

The baseline parameter calibration and the choice of initial conditions are summarised in Table 1. We choose one day as the time unit in all simulations.¹⁵

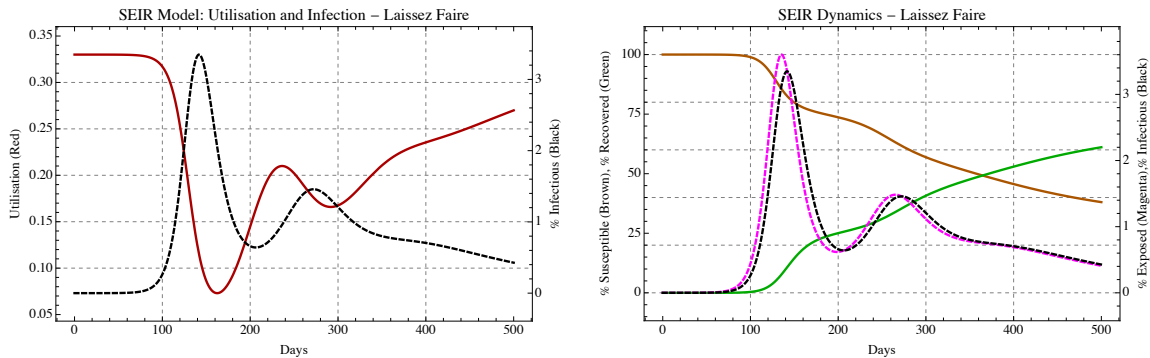
Parameter / Initial Condition	Value	Moment to match	Source
β_0	1.27	$r(0) = 3$	Flaxman et al. (2020); Li et al. (2020)
γ	0.143	$\gamma^{-1} = 7$	Prem et al. (2020)
δ	0.2	n/a	Authors' calculation
α	0.025	n/a	Authors' calculation
σ	0.143	$\sigma^{-1} = 7$	Atkeson et al. (2020)
$u(0)$	0.33	$\bar{u} = 0.33$	Piketty (2014)
$E(0)$	$0.35 * 10^{-4}$	n/a	Authors' calculation

Table 1: Baseline parameter calibration

3.1 Laissez-faire

Our first exercise is to examine the joint dynamics of the epidemiological and economic variables in a *laissez-faire* scenario in which the pandemic runs its course, and the government adopts no measures. The left panel of Figure 2 illustrates the dynamics of the main two variables – the utilisation rate and the infection rate – against time. The right panel focuses on the epidemiological variables.

Figure 2: SEIR- u model: laissez-faire



The dynamics of the integrated SEIR- u model is rather different from that of the standard SEIR model, in which the interaction between economic activity and the epidemic is ignored. In particular, rather than having a monotonic increase in infections, followed by a monotonic decline once

¹⁴See Appendix B for more details on the calibration of the model.

¹⁵The datasets used below, as well as the *Mathematica* code used for these simulations – including a variety of additional robustness checks and extensions – is available at <https://www.danieletavani.com/code>.

herd immunity is reached, the SEIR- u model predicts a cyclical movement in *both* epidemiological *and* economic variables, with successive waves of infection interacting with business cycles.

At the onset of the epidemic the economy is running at full capacity. As the virus is left free to spread, the epidemic hits the workforce and economic activity drastically declines. The decline in the utilisation rate yields a reduction in contacts, which in turn leads to a drop in infections, just as in the 2D model. As economic activity resumes, so does the growth rate of the pandemic, and the cycle repeats, although with lower intensity, leading the economy to move, after about a year and a half, in the vicinity of the steady state without infections. The fact that the endemic steady state – which was the crucial feature of the 2D model – disappears in the full SEIR model is due to the economy eventually crossing the herd immunity threshold¹⁶

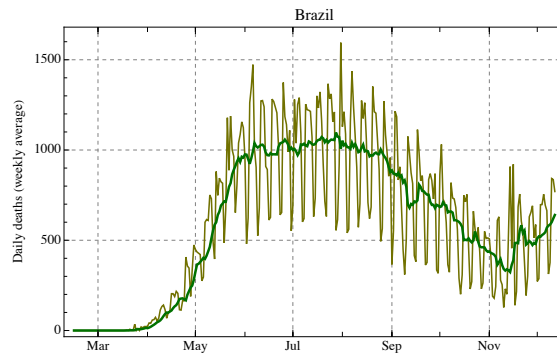


Figure 3: Daily deaths and rolling 7-day average (green) in Brazil, February 15th–December 12th, 2020. Source: Worldometer.

Does the cyclical dynamics of the SEIR- u model capture a relevant mechanism of the joint interaction of economic and epidemiological variables in the real world? Figure 3 displays data on daily deaths for SARS-CoV-2 in Brazil,¹⁷ whose Federal Government has purposely elected not to implement any stringent regulations on behaviour and economic activity, and sometimes even actively undermined the policies adopted by independent bodies and local governments.¹⁸

Perhaps even more informative is the comparison of different US states. For there exists “significant variation across states . . . For example, Republican-led and Midwestern states have the lowest levels of policy stringency even as they face increasing case counts. These trends hold particular

¹⁶This happens around 300 days into the epidemic in Figure 2 due to the chosen parameter calibration, which draws on the consensus in the literature. A lower $r(0)$ (as long as it is above 1) would result in the epidemic lasting longer, and consequently in a longer period of economic fluctuations.

¹⁷Given the massive changes in testing capacity and policy since the beginning of the pandemic, data on daily new cases are rather unreliable. Although factors such as the response and ability to cope of health systems has also changed over time, data on daily new deaths are likely to better capture broad trends in infections. It is worth keeping in mind, however, that we are not trying to validate our model econometrically, and any comparisons between our model and individual countries are merely illustrative.

¹⁸The stringency of government measures adopted worldwide to tackle the pandemic has been measured on a daily basis by Hale et al. (2020b). Although Brazil has had a relatively high stringency index for most of the past ten months this is due to the uncoordinated and largely ineffective actions taken by state and local governments (for a comprehensive discussion, see Petherick et al. (2020)). Besides, if one normalises the stringency index dividing it by the number of cases per million inhabitants, Brazil’s average stringency index during the last several months is among the lowest in the G20.

relevance in the absence of a centralized, federal response” (Hale et al., 2020a, p.3). Figure 4 shows data on daily deaths for SARS-CoV-2 in the US States with the lowest stringency indices.¹⁹

To be sure, no stylised low-dimensional model pitched at a high level of abstraction can provide an accurate depiction of social reality. Yet the pattern depicted in Figure 2 is consistent with the empirical evidence of those states who have adopted a relatively lax approach (Figures 3 and 4). Indeed, with the exception of states which seem to be currently facing the first wave of the epidemic, the similarity with the cycles depicted in Figure 2 is striking. This suggests that the SEIR-*u* model does capture *some* important mechanisms linking economic activity and the dynamics of infections.

Three features of the dynamics displayed in Figure 2 should be emphasised, which are relevant for the welfare implications of the pandemic. First, absent any government intervention, the significant increase in infection rates in the first wave is likely to lead to major welfare losses, including deaths, as the health system is overrun (see section 3.3 below). Second, the economy experiences a sharp downturn, and wide fluctuations in economic activity. Third, the disruption is not short-lived: in the simulation presented, it takes well over a year and a half for the epidemic to be under control with economic activity still below the normal level.

In summary, based on the available evidence on the infectiousness of SARS-CoV-2, and on rather conservative estimates of the effect of the epidemic on economic activity, our results cast significant doubts on laissez-faire strategies. The virus does not disappear on its own within a reasonable time frame, and the economy does not quickly return to normal, possibly after some turbulence. What are, then, the most appropriate measures to tackle it? To this question we turn next.

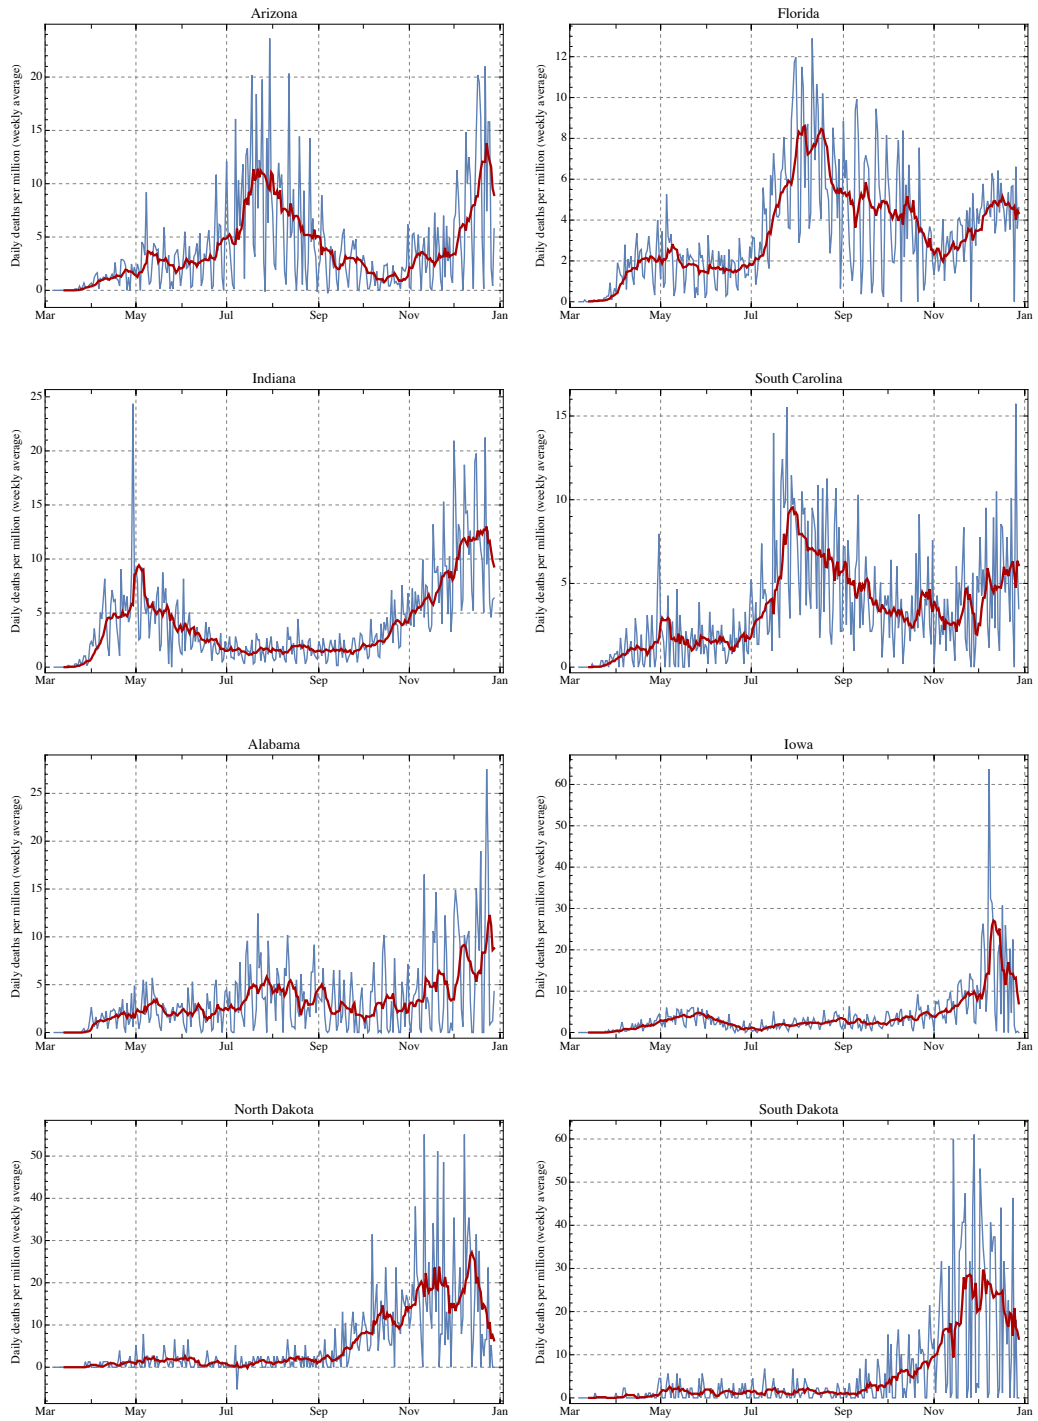
3.2 Policy

One of the key features of the pandemic has been the very wide range of measures adopted globally: from hard lockdowns to mild physical distancing, to almost uncoordinated business-as-usual. The wide variety of policies adopted in different social and economic contexts is one of the reasons why the pandemic may be considered as a unique natural experiment in comparative political economy.

It is impossible to incorporate the vast range of measures adopted by any one government during the pandemic in a single formal model, and indeed it would not even be particularly informative to do so, as it would not allow us to disentangle their effects. Our aim here is different: we use our analytical framework in order to examine alternative policies *in vitro*. We focus on what may be considered the two main approaches adopted in the last few months – namely, physical distancing measures and hard lockdowns. First, we analyse them separately, then we consider a more complex policy mix in which restrictions on economic activity interact with distancing advice.

¹⁹We normalise the data focusing on deaths per million inhabitants in order to facilitate comparisons with US states with a high stringency index in Figure 9 below.

Figure 4: Daily deaths per million inhabitants, and rolling 7-day average (red) in US States with low stringency index, March 1st–December 28th, 2020. Source: The COVID Tracking Project at *the Atlantic*.



3.2.1 Physical distancing measures

One model policy approach to the epidemic consists in enforcing physical distancing regulations and advising the public to refrain from certain social activities, without restricting economic activity. The main effect of physical distancing, whether enforced by government policies (e.g. with regulations concerning the use of public, or even private, spaces) or resulting from behavioural changes, is to reduce the number of contacts, and thus the reproduction rate $r(t)$. We can capture this effect by assuming that physical distancing has a scaling effect on the infection probability.

Formally, let $\lambda \in [0, 1]$ denote the proportion of physical contacts occurring after the introduction of distancing prescriptions: the parameter captures both the stringency of policy and the behavioural response of the public. If no restrictions on social activities are imposed and agents do not change their behaviour, then $\lambda = 1$ and the number of contacts remains unchanged. Values of λ below 1 correspond to increasingly stringent distancing advice, and more radical behavioural changes, leading to a greater reduction in contacts. Then, equations (6)-(7) can be generalised as follows:

$$\dot{S}(t) = -\frac{\lambda\beta_0 u(t)S(t)I(t)}{N}, \quad (16)$$

$$\dot{E}(t) = \frac{\lambda\beta_0 u(t)S(t)I(t)}{N} - \sigma E(t). \quad (17)$$

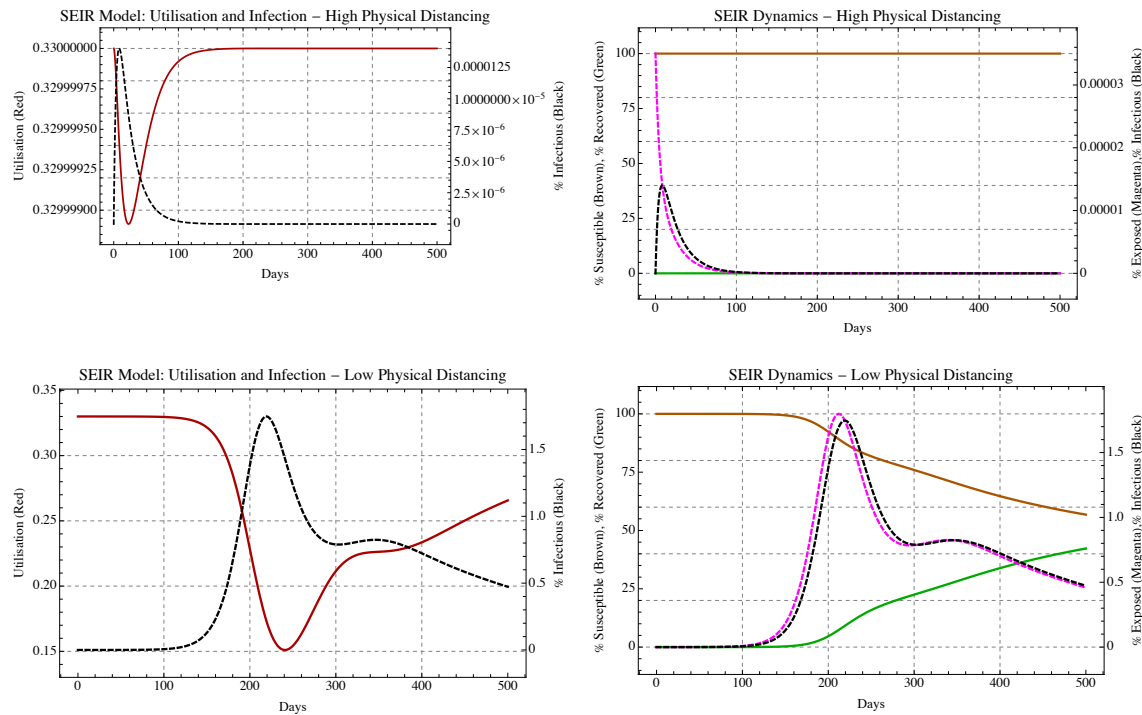
Given the exogenous determinants of the probability of infection β_0 , the rate of utilisation $u(t)$, and the recovery rate γ , equations (16)-(17) seem to suggest that health authorities can target the amount of physical distancing, λ , required to significantly mitigate the effect of the virus and in the limit even ensure that the disease die out by enforcing distancing measures such that $\beta_0\lambda u(t) < \gamma$. This seems to capture the beliefs of governments, such as, famously, the Swedish one, who have elected to avoid harder measures, while enforcing basic physical distancing measures, and encouraging the public to reduce social activities. In so doing, the government hoped to reduce the impact of the virus without major economic losses.

If the joint action of government regulation and changes in behaviour and social norms were able to reduce contacts significantly (yielding a sufficiently low λ), then the epidemic would die out within a few months while causing a mild and short-lived adverse economic effect. Such scenario is represented in the top panels of Figure 5 corresponding to a distancing parameter $\lambda = .15$.

Yet, no society can reach the levels of reduction in social interactions necessary for the epidemic to stop without affecting economic activity. One can use the value of the reproduction rate at the beginning of the pandemic, $r(0) = 3$, and the condition that $r(0)\lambda < 1$ for the disease to die out to find that λ must be smaller than $1/3$ in order for physical distancing to contain the epidemic without restricting economic activity.²⁰ It is completely unrealistic to think that physical distancing alone can reduce contacts by around two thirds (let alone 85%) in the entire population without any policy measures or behavioural reactions directly or indirectly affecting the economy.

²⁰It is not difficult to show that $\lambda = 1/3$ is the cut-off value below which the no infection steady state of the simplified model in section 2.3 switches from locally unstable to locally stable.

Figure 5: SEIR- u model: strong vs. weak physical distancing



How about lower levels of physical distancing? The right panel of Figure 5 depicts the behaviour of the model with a constant 30% reduction in contacts ($\lambda = .7$).²¹ It is immediately apparent that weaker physical distancing yields an evolution of the epidemic that is qualitatively identical to the laissez-faire case: health advice and the related behavioural response may lead to a lower peak in infections as well as a lower drop in economic activity, but both reductions are relatively small.

The limited effect of physical distancing is all the more remarkable if one notes that, while less far-fetched, a 30% across-the-board reduction in contacts underlying the right panel of Figure 5 remains extremely unlikely. As shown by the Swedish experience, and contrary to the hopes of the country’s government, it is not possible to enforce high enough physical distancing without any additional measures affecting economic activity – in our framework, the rate of capacity utilisation u . There are only a limited number of social activities whose reduction has no impact on economic activity, and there is only so much one can do by, say, reducing entry in public playgrounds and parks, requiring people to avoid family gatherings, or limiting the number of social visits, to mention a few. In other words, it is impossible to reduce contacts to sufficiently low levels – i.e. to achieve a sufficiently low λ – without interfering – directly or indirectly – with economic activity.

Figure 6 shows data on daily deaths in Sweden: the pattern is remarkably similar to that in the

²¹ Kaplan et al. (2020) also distinguish reductions in contacts due to decreases in economic activity and an exogenous component capturing “reductions in the transmission rate ... that arise from changes in social norms about cheap preventive measures, such as wearing masks, washing hands” (Kaplan et al., 2020, p.18). They consider reductions in such exogenous component of 20%. Other empirical papers argue for larger reductions in the order of 40% (see, for example, Mitze et al., 2020). We have chosen an intermediate value.

lower panels of Figure 5 and it clearly illustrates the limitations of physical distancing policies.

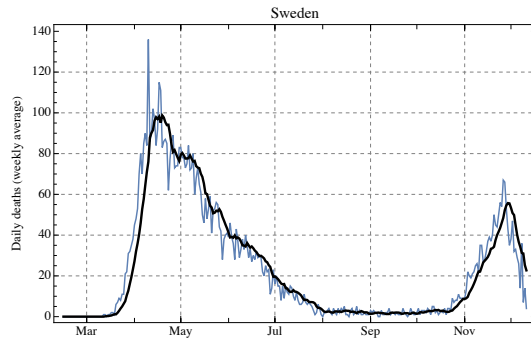


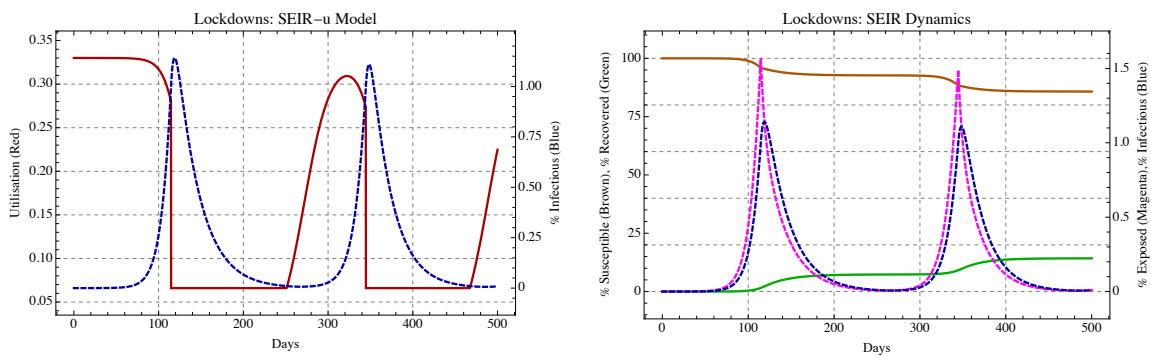
Figure 6: Daily deaths and rolling 7-day average (black) in Sweden, February 15th–December 12th, 2020. Source: Worldometer.

3.2.2 Lockdowns

One of the most widespread approaches adopted throughout the globe at the peak of the pandemic has been to impose direct restrictions on economic activity. These restrictions have ranged widely in terms of their severity, and the economic sectors involved, but the common feature of all lockdown measures has been a reduction of economic activity relative to potential – namely, a reduction in capacity utilisation. Thus, our framework provides a natural way of modelling lockdowns as hard caps on utilisation, u , imposed whenever certain conditions obtain.

In the simulation displayed in Figure 7, we assume that when the infection rate reaches the threshold value of $\bar{I} = 1\%$, the government puts in place measures that shut down activity in a large part of the economy, reducing u to 20% of its normal value and keeps them in place until the infection rate goes below $\underline{I} = .01\%$.²²

Figure 7: SEIR- u model: lockdowns



²²This choice of parameters and the focus on the infection rate – rather than, say, the reproduction rate – is merely illustrative: as we already noted, lockdown policies have varied greatly across time and space. Our main results are robust to a large set of perturbations of the specification of lockdown policies (see section 5.3 below).

Figure 7 clearly shows that lockdowns provide an extremely effective way of reining the infection in. When a lockdown occurs, economic activity immediately drops to the cap, which curbs the growth rate of infections and the peak of the first wave is significantly lower than in the laissez-faire scenario. Yet, as the reduction in contacts affects the spread of the epidemic only with a lag, the lockdown can only be lifted after a certain period. The slump is large, and the loss of output significant, but it is also rather short and the infection rate drops rather quickly. Indeed, a comparison of Figures 2 and 7 immediately shows that, given the independent negative effect of the epidemic on economic activity, the main effect of a lockdown is to make the downturn only slightly more pronounced, but significantly shorter compared with the laissez-faire.

Further, even when the lockdown is lifted, and economic activity resumes, – as it happened throughout the world at the end of the Spring of 2020 – the infection rate remains very low for a rather long period – in our model, about four months, which is similar to the experience of most European countries after the first wave. Of course, following the upturn in economic activity, and the consequent surge in contacts, the pandemic dynamics eventually resumes. But the main goal of lockdown policies is not to eliminate the virus: it is to buy time, dramatically reducing infections – albeit at the expense of sharp but temporary downfalls in economic activity – while the population (and the economy) waits for widespread vaccination.

Keeping in mind the caveat that our simple model does not aim to provide a comprehensive, detailed picture of any given economy, a quick look at the data of countries and US states which have adopted stringent restrictions on economic activity suggests that the SEIR- u model does capture some relevant aspects of lockdowns. Figures 8 and 9 display the dynamics of daily deaths, respectively, in France and Spain – two countries that have imposed strict lockdown policies at the various stages of the epidemic – and in select US states among those with the highest stringency index (Hale et al., 2020a). Again, the broad similarity with the simulation results is hard to miss.

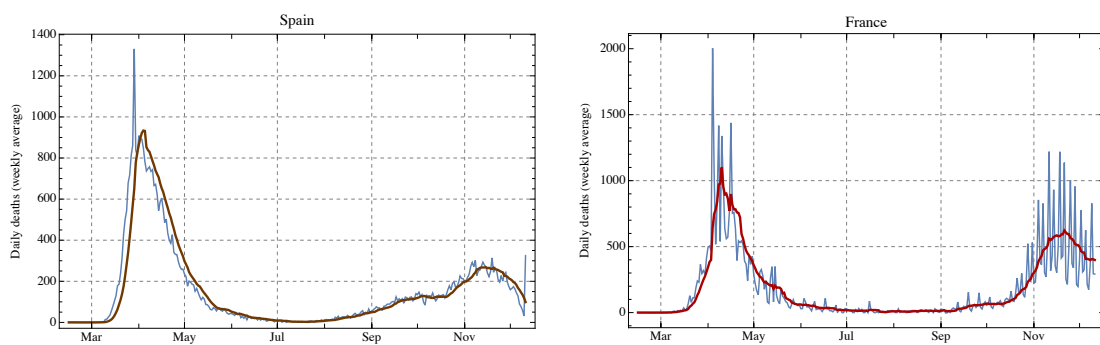
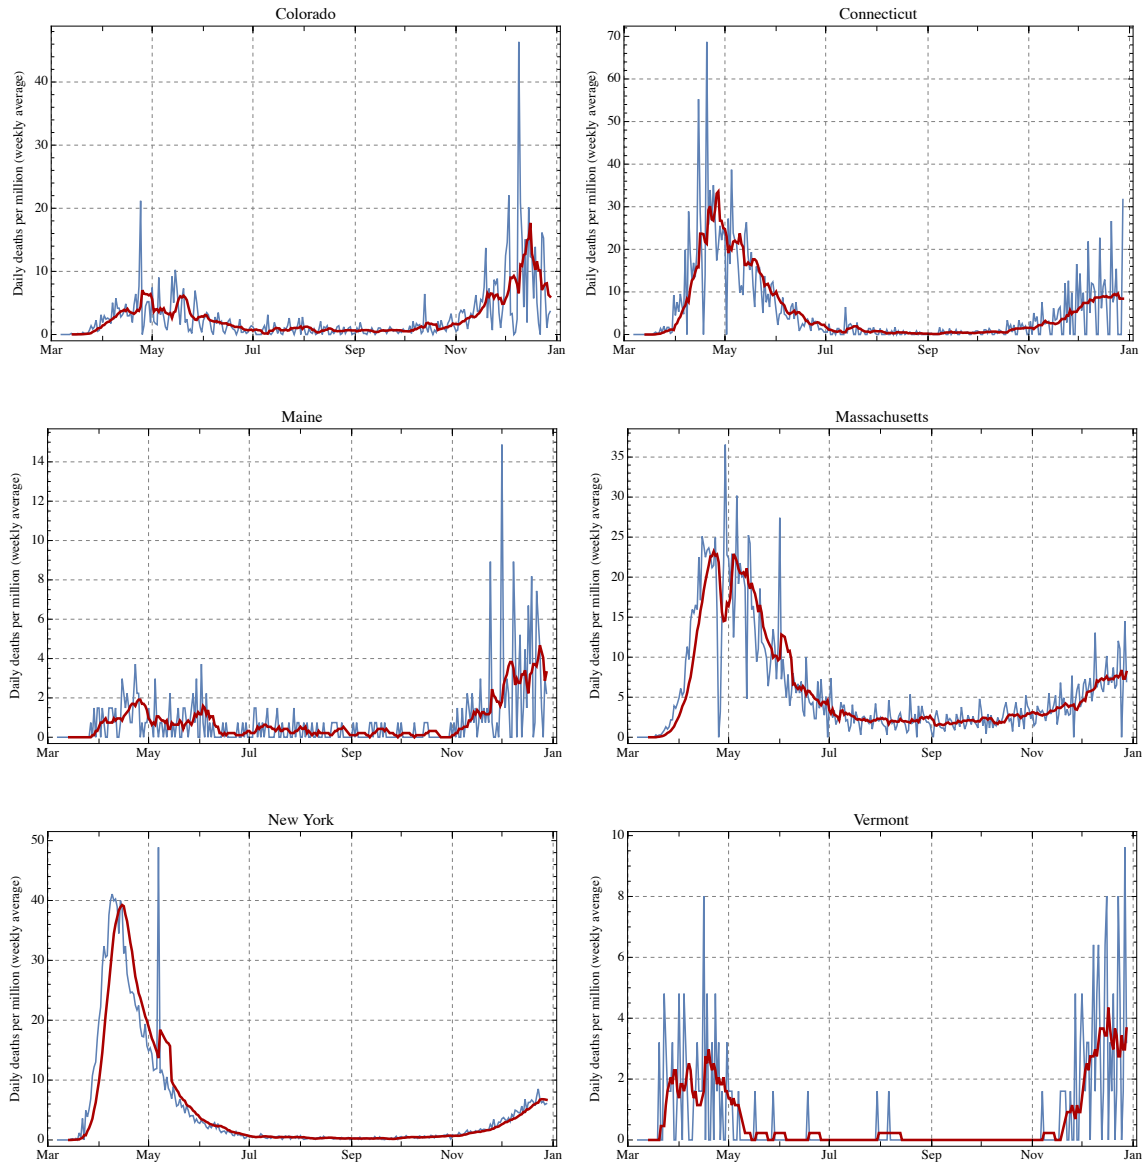


Figure 8: Daily deaths and rolling 7-day average (red) in Spain and France, February 15th—December 12th, 2020. Source: Worldometer.

Figure 9: Daily deaths per million inhabitants and rolling 7-day average (red) in US States with high stringency index, March 1st–December 28th, 2020. Source: The COVID Tracking Project at *The Atlantic*.



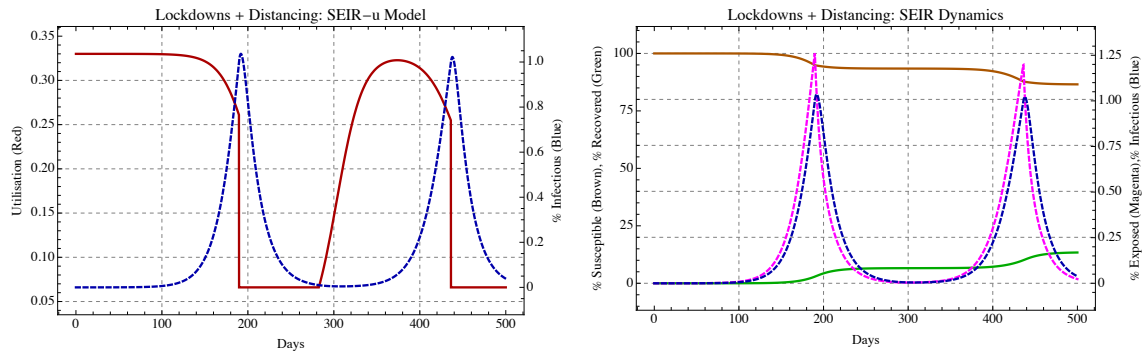
3.2.3 Lockdowns/Physical distancing interactions

Even though it is informative to look at distancing policies and lockdowns separately, and the two sets of measures can and should be separated conceptually, in reality they tend to interact, and lockdowns have been imposed in conjunction with distancing recommendations. As we noted in section [3.2.1](#), it is not possible to promote sufficient physical distancing without affecting the economy. Conversely, physical distancing measures are more effective when restrictions on economic activity are also in place. Relatively high levels of physical distancing are likely attainable when schools and

higher education institutions are closed, occupancy limits are imposed in the service and hospitality sector, as well as theatres, concert halls, etc. Further, part of the effect of lockdowns is not just in terms of reduction of contacts on the workplace, but also on work-related social interactions.

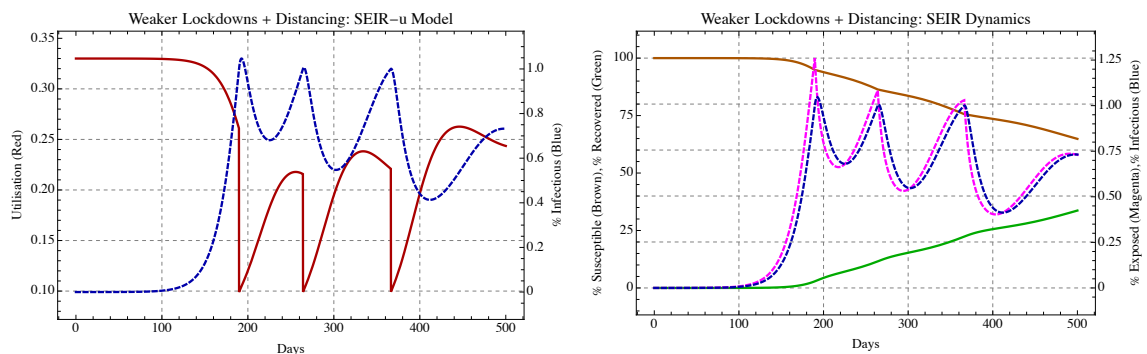
It is straightforward to integrate the two policies, as displayed in Figure 10 below. For this simulation, we use $\lambda = .7$ corresponding to a reduction in contacts among the population of 30%, and the same specification of lockdowns as above.

Figure 10: Lockdowns/physical distancing interactions.



An alternative policy experiment would involve the government authority setting a less stringent cap on economic activity, counting on compliance by the public on physical distancing recommendations. Figure 11 displays a simulation where the cap on economic activity is set to 30% of normal, instead of 20%. The trajectory followed by the economy appears to conform with common sense: while in principle there is a tradeoff between the stringency of a lockdown and the behavioral reduction in contacts by the public, relying extensively on the latter – while relaxing the former – exposes a society to higher epidemiological risks, while sudden and frequent stoppages are likely to be more disruptive than longer but sporadic interruptions of economic activity.

Figure 11: Less stringent lockdowns with physical distancing.



3.3 Death rates

A proper welfare analysis would require the explicit formulation of a social welfare function and a comprehensive description of the multifaceted effects that different policies have throughout the economy. This goes well beyond the scope of our simple model. Nonetheless, one key objective of *any* public policy tackling the effects of the epidemic is to minimise deaths.

One major advantage of lockdown policies is precisely that they drastically reduce the number of deaths associated with an epidemic. Using data on death rates from the United States we can calculate the case-fatality rate of the disease as the ratio of total deaths over total reported cases. On December 29th, 2020, the New York Times reports a total of 19.3 million cases and 335,141 deaths. The left panel of Figure 12 uses the corresponding ratio to plot the number of individuals who die in different scenarios in the SEIR- u model. Two features of the simulations immediately stand out. First, although there is no outright dominance throughout, it seems clear that uncoordinated, laissez-faire policies are the worst possible approach from the standpoint of public health.

Second, the simulations clearly show that lockdowns should be a fundamental ingredient of any approach to the epidemic. To be sure, it may seem difficult to rank a pure lockdown policy and physical distancing measures, as the former appears to be more effective only in the later stages of the epidemic. This is only because, in order to analyse different policies *in vitro*, we have assumed that until the imposition of the lockdown, the government is passive: lockdown measures are effective *once they are introduced* but in the meantime a laissez-faire approach is *de facto* adopted. In contrast, physical distancing measures are adopted from the very beginning of the epidemic.

If one considers the more realistic policy mix discussed in section 3.2.3, in which the government imposes physical distancing measures immediately while being prepared to impose severe restrictions on economic activity, then the importance of lockdown policies from the standpoint of public health is evident.²³

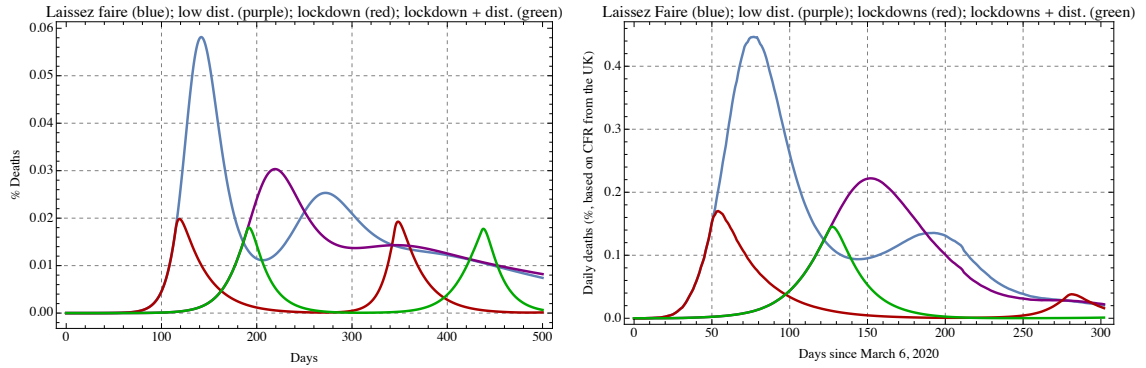
These results are all the more remarkable given that we are using a constant case-death ratio at all stages of the epidemic, and regardless of the number of infected individuals. This is equivalent, *de facto*, to assuming that there exists no capacity constraint in the health sector – a strong assumption during a pandemic. If one assumed that the case-death ratio increases once infections cross a certain threshold, then the advantage of lockdown policies – whose main effect is precisely to prevent any spikes in infections, at the possible cost of a slightly longer diffusion of the disease – would be even more evident. As an illustration, the right panel of Figure 12 multiplies the numerical solution of infection rates corresponding to different policy scenario (laissez faire, physical distancing, lockdowns, and lockdowns with physical distancing) by the *observed* case-fatality rate in the United Kingdom since March 6th, 2020, which has indeed spiked during the first wave, and then slowly declined thereafter and until the second wave.²⁴ This thought experiment highlights

²³The effectiveness of lockdown policies in reducing deaths is even clearer in the simplified, short-run 2D model; see the online addendum at danieletavani.com/code. Of course, less stringent lockdown policies, such as those considered at the end of section 3.2.3, yield less dramatic reductions in deaths.

²⁴Naturally, the actual case-fatality-rate is endogenous to the adoption of mitigation policies by the government. However, because infections at time t in all the various policy scenarios are multiplied by *the same* CFR at time t – just like in

even more starkly the superiority of a combined lockdown - physical distancing policy mix.

Figure 12: Death rates: constant case-fatality rate (left) vs. variable case-fatality rate (right). Sources: The New York Times and Our World in Data, University of Oxford.



4 Distributional dynamics

From a policy perspective, controlling the epidemic and mitigating the fall in output are certainly focal concerns, especially in the short run. Nonetheless, as we noted in the Introduction, losses in aggregate income are only one of the economic consequences of the pandemic. Equally relevant are its distributional implications, and in particular the effect on workers' incomes: having a smaller set of alternatives, employees are likely to bear the brunt of the epidemic.²⁵

We examine the distributive implications of the epidemic by adding a dynamic equation to incorporate the evolution of income shares in the model. We focus on the wage share ω , defined as the ratio of real wages to labour productivity. In standard structuralist and post-Keynesian models, its rate of change is typically assumed to increase in economic activity as the labour market tightens and wages rise relative to labour productivity. In order to obtain an interior steady state, and because the wage share is bound above by 1, we also assume a stabilizing mechanism so that further increases in the wage share reduce its rate of change. Finally, we incorporate a “drift” parameter $\eta > 0$, which is necessary in order for the wage share not to drop to unreasonably low values when utilisation becomes very small. The equation describing the evolution of the wage share is²⁶

$$\dot{\omega}(t) = [\eta + \theta u(t) - \phi \omega(t)] \omega(t), \quad (18)$$

where all parameters are assumed to be positive and numerically calibrated in order to obtain a value

the left panel of the figure – the figure is informative because it allows for a fair comparison of alternative measures.

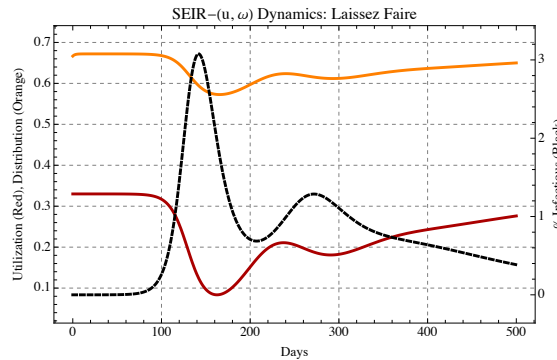
²⁵Self-employed people are also disproportionately affected by the reduction in economic activity. It is now customary in most datasets to adjust the measure of the wage share for self-employed income: our stylised model of distributional dynamics should be understood as including self-employed income.

²⁶Considering the simplified 2D model in section 2.3, the long-run values for the labour share corresponding to the no-infection steady state and the endemic steady state are obtained as: $\omega^{NI} = \frac{\eta + \theta \bar{u}}{\phi}$ and $\omega^E = \frac{\eta + \theta u^E}{\phi}$. Given our assumptions, the endemic equilibrium is locally asymptotically stable, while the no infection steady state is unstable.

of the wage share roughly equal to two-thirds of national income when the economy is at capacity.²⁷

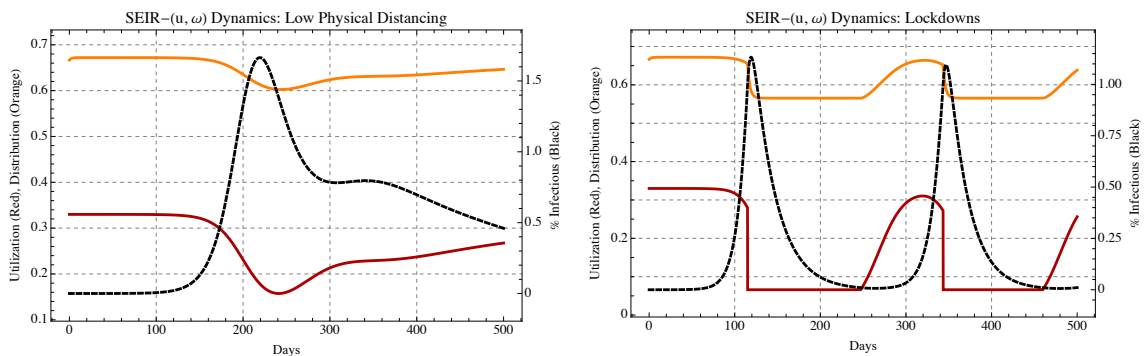
Figure 13 shows the simulation results for a country that adopts a laissez-faire policy and lets the infection run its course. The epidemic causes a significant decrease in the wage share (about 7-8 percentage points in our simulation) which does not return to its long run value for a very long time. This arguably provides a further reason for the government to intervene.

Figure 13: SEIR- (u, ω) model: laissez-faire



Do policy measures aimed at tackling the epidemic also moderate its distributional effects?²⁸ Figure 14 shows the dynamics of the virus, economic activity, and, through the latter, income distribution, corresponding both to a physical distancing policy with $\lambda = .7$ and to a hard lockdown. Two features are immediately apparent. First, even if a 30% reduction in contacts was indeed feasible without intervening in the economy, it is clear that besides being relatively ineffective in dealing with the epidemic, this would have a rather minor effect on the fall in the wage share.

Figure 14: SEIR- (u, ω) model: physical distancing vs. lockdowns



Second, as argued in section 3.2.2, lockdown policies are way more effective in dealing with the epidemic at a comparatively limited cost in terms of output loss. Yet, their distributional effects are relatively limited. Caps on economic activity as infections spike lower real wages given the

²⁷Specifically, we assume $\eta = 1/3$, $\theta = .25$, $\phi = .68175$. All other parameters are as specified in Table 1.

²⁸In the simplified model in section 2.3, any policy that gets the epidemic completely under control, and avoids the endemic scenario, would be ultimately beneficial to income distribution, as $\omega^E < \omega^{NI}$ (see footnote 25).

resulting slackness in the labour market. As infection subsides and lockdown measures are lifted, the wage share rises: but inevitably, the infection rate spikes again, which forces another lockdown and consequently a reduction in the wage share.

Although in both scenarios the wage share eventually returns to its normal value in the long run, a major slump, or wide fluctuations in the wage share over a prolonged period of time are arguably undesirable. Moreover lockdown measures do tend to produce larger falls in the wage share compared with the laissez-faire scenario.

This is an important point, both from the viewpoint of distributive justice and in terms of policy feasibility. For, distributing more evenly the burden of public health policies is not only fair: it is an important ingredient to insure compliance, and therefore ultimately the success of any policy measures.²⁹ This provides a robust justification for public intervention aimed at attenuating the economic effects of the epidemic, *and* of the measures to tackle it, for important segments of the population. Actually, although various forms of subsidies may be introduced (or increased) in order to make up for any income losses, our model suggests that employment protection policies, and bans on layoffs are equally, if not more important.

We therefore examine some simulations in which labour-market policies such as unemployment insurance or direct transfers are introduced in order to partially shield workers from the loss of wage income arising from layoffs occurring during epidemic-induced recessions.

There is no obvious way of introducing active labour market policies in our model. Theoretically, in the structuralist literature, labour market conditions affect distribution via the bargaining power of the working class. It is therefore natural to assume that if the government wants to mitigate the adverse distributive effect of the epidemic, it will intervene with policies that weaken the link between economic activity, labour market conditions, and wages.

This intuition could be formalised, for instance, by turning the parameter θ , which captures the cyclical dependence of the wage share on utilisation, into an inverse function of economic activity, so that $\theta = \theta(u)$ with derivative $\theta_u < 0$. In the simulation displayed in Figure 15, we used a specification $\theta(u) = \theta_0 u^{-p}$, $\theta_0 = .05$, $p = 1.1$ which implies a mildly countercyclical response of the wage share to economic activity through the effects of unemployment insurance or direct transfers to households. It is worth emphasising that this simulation aims to illustrate the importance of active labour market policies, and it is not meant to actually replicate any specific measures adopted by a given country or local government.

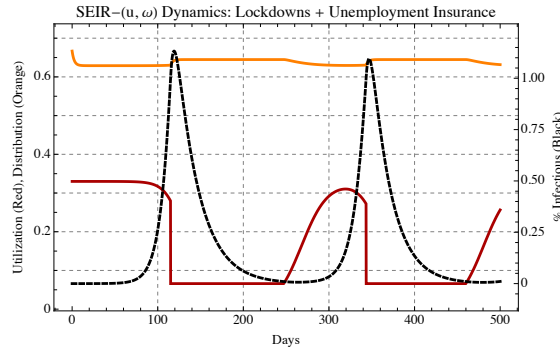
5 Robustness checks and extensions

We have analysed many variations of our models in order to assess the robustness of our results. This section briefly summarises the main points.³⁰

²⁹In this paper, we ignore the issue of compliance by implicitly assuming all policies to be automatically enforced. This is merely for analytical purposes as it allows us to compare all measures in vitro. However, it is rather unclear whether different policies are systematically associated with different levels of compliance.

³⁰More details can be found in the online Addendum at www.danieletavani.com/code.

Figure 15: Lockdowns with unemployment insurance.



5.1 Alternative calibrations

In order to capture a situation corresponding to the very beginning of an epidemic – before it has any perceivable effect on the economy – all of our simulations have been initialised assuming that the economy is at the normal level and only a very small proportion of the population is infected. Different values of $E(0)$ slightly modify the shape of the first wave of infections (higher values imply a slightly higher first peak) but nothing else changes.

Similarly, a normal rate of utilisation equal to one third is common in advanced countries, however the model predictions are virtually unchanged for a wide range of values of \bar{u} , up to .8 if one interprets capacity utilisation as done for instance by the Federal Reserve in its business survey. The value of \bar{u} affects the endemic infection rate in the 2D model; but it is of no consequence on the qualitative features of the transitional dynamics, either in the 2D model or in the full SEIR- u .

Estimates of the duration of the infectiousness of SARS-CoV-2 vary greatly depending on the type of cases analysed (symptomatic vs. asymptomatic), the relevant demographics (age groups, gender), and the methodology adopted in the study. We have set $\gamma = 1/7$, reflecting an average duration of infectiousness of 7 days based on meta-analyses of the clinical literature, but we have run our simulations allowing $1/\gamma$ to range from 5 to 10 days. Calibrating β_0 to match $r(0) = 3$, the qualitative behavior of the model is not affected by the chosen value of the parameter γ .³¹

We have analysed the behaviour of the model for a range of values of $r(0)$ between 2.5 and 4. Lower values of the simple reproduction rate imply a slower diffusion of the virus, which delays the peak of the epidemic and also the crossing of the herd immunity threshold. But other than that, nothing changes. Indeed, the qualitative dynamics of the model are unchanged provided $r(0) > 1$, or $\beta_0\bar{u} - \gamma > 0$. If $\beta_0\bar{u} - \gamma < 0$, then starting from a normal utilisation rate and a low level of the infection rate, the infection causes a barely noticeable, temporary drop in economic activity and, as the infection dies out, the utilisation rate converges monotonically back to its normal value. The simulation corresponding to such a scenario – which could describe a regular flu season for example – mimics very closely the path under high physical distancing in Figure 5.

Because there is no natural way of calibrating the speed of adjustment parameters α and δ , in

³¹In the 2D model, at the endemic equilibrium the infection rate varies with γ , while the utilisation rate is unaffected.

our baseline simulations, we have adopted a conservative strategy choosing values that do not drive the dynamics of the model. However, we have run a number of robustness checks for lower values of δ as well as higher values of α . The behaviour of the model does not change.

Finally, the analysis of the interplay between the pandemic and distribution in section 4 are robust to a wide variety of perturbations of the parameters η , θ , and ϕ provided they are calibrated to constrain the wage share to remain in the $[0, 1]$ interval. Changes in these parameters affect the amplitude of the boom-bust cycles in laissez-faire, but not the qualitative behavior of the model.

5.2 Alternative model specifications

In addition to the main SEIR- u model, we have analysed in depth the simplified 2D model in section 2.3 in order to focus on the short-run effects of different policies. The results are even starker: realistic levels of physical distancing make hardly any difference either to the dynamics of the infection, or to economic activity. The economy displays cyclical dynamics very similar to those of the laissez-faire scenario around a marginally-improved endemic steady state. Instead, lockdowns act swiftly and effectively to reduce infections although at the cost of slightly bigger (cyclical) reduction in economic activity.

All of our conclusions remain unchanged if we obliterate the distinction between exposed and infectious individuals, and simplify the epidemiological bloc of our model into the so-called SIR model. The main difference with respect to our core SEIR- u framework is that lockdown policies have instantaneous effects and start curbing the contagion as soon as they are imposed.

Because $\beta(t)$ captures a probability, the multiplicative structure in equation (5) is a natural way of modelling the effect of economic activity on the dynamics of the epidemic. Nonetheless, we have examined the model under alternative assumptions. Specifically, we have considered the family of exponential functions $\beta(t) = \beta_0 u(t)^\sigma$, $\sigma \in [1, 2)$, of which the linear model is a special case. This makes the predicted swings in economic activity slightly more pronounced, but as long as σ is below 2 – which seems quite plausible given the observed interaction between economic activity and infection – the dynamics of the model remains essentially unchanged.

There is no obvious way of choosing the function $\theta = \theta(u)$ describing the effect of active labour market policies in the model presented in section 4. We have opted for a functional form reflecting rather conservative assumptions but our conclusions are robust both to a wide range of choices of the parameters θ_0, p and to a number of different functional forms (with derivative $\theta_u < 0$). For instance, higher values of the policy parameter p result in more countercyclical behaviour of the wage share; linear functional forms have similar implications.

5.3 Alternative policy specifications

Concerning physical distancing, we ran a number of simulations with alternative values of λ . For any value of λ below the minimum required physical distancing – which, importantly, varies with β_0 – the pandemic dies out on its own; while for any value of λ above the minimum required physical

distancing, the model behaves similarly to the laissez-faire scenario.

We have also run simulations corresponding to lower as well as higher mandated reductions in economic activity. In the former case, lockdowns tend to be shorter and more consequential in terms of reductions in output; in the latter case, the crucial cutoff for the lockdowns to bite is to impose a lockdown rate of capacity utilisation below the endemic utilisation rate in the 2D model, otherwise the behaviour of the economy will revert to the laissez-faire case.

Similarly, we have run simulations with lower and higher values of the threshold values for the activation and lifting of lockdown measures, respectively \bar{I} and \underline{I} . If \bar{I} is lowered, economic activity grinds to a halt sooner leading to a more significant impact on infections, and a more significant reduction in deaths, but the measures need to stay in place for a shorter period.

We have also explored different ways of conceiving of lockdown policies. To be specific, first, we have analysed a different specification with lockdowns conceived of as sharp but temporary drops in \dot{u} . Compared with laissez-faire (or realistic physical distancing measures), stop-and-go measures of this sort produce sharper decreases in infections, but their effect is short-lived: disruptions in economic activity are therefore temporary but recurring and frequent.

Second, we have run simulations assuming that hard caps on utilisation are imposed after the *reproduction rate* – rather than the infection rate – crosses a certain threshold; specifically, one. This is justified given the concern of policymakers with hospital capacity and with the dismal arithmetic of COVID-related deaths rising about two weeks after spikes in the reproduction rate. Lockdowns of this kind deliver low infection rates for a longer period by more pronounced shutdowns in economic activity. Further, the bounce-back is slower to unfold, and subsequent swings in economic activity (and infection) are less tamed than when lockdowns are conditional on the infection rate.

Third, we have checked the robustness of the results concerning the deaths associated with alternative policy approaches when the case-fatality rate is allowed to change over time, as in the right panel of Figure [12](#). Our conclusions remain unchanged if, instead of the actual case-fatality rate of the United Kingdom, one focuses on any of the countries that we have considered as examples of the different approaches to the pandemic, such as Brazil, the US, Spain, France, or Sweden.

Finally, as we argued in section [3.2.3](#), the level of physical distancing that can be reached in a given society depends on economic activity. Lower levels of economic activity arguably make physical distancing measures both easier to comply with and more effective. Formally, $\lambda = f(u)$, and we have simulated our model postulating a simple affine form for f as follows:

$$\lambda(t) = \lambda_0 + \lambda_1 u(t), \quad (19)$$

where λ_0 captures exogenous factors such as health advice and behavioural changes (including wearing masks and other cheap preventive measures), while λ_1 captures the interaction between the utilisation rate – our proxy for economic activity – and non-work-related physical distancing.

We have run various simulations using different values of λ_0, λ_1 .³² Given the nonlinear effect

³²In the online Addendum we present the results of two simulations in which we calibrated $\lambda(t)$ using mobility data for the UK from [Di Guilmi et al. \(2020\)](#), in conjunction with the parameterisation in Table [1](#).

of utilisation on the probability of contacts, these simulations deliver shorter lockdowns (lasting about three weeks) followed by about three months of infections close to zero before the infection rate climbs up again to the threshold required for a new lockdown to take effect. Otherwise, the dynamics are very similar to the mixed policy scenario in section [3.2.3](#).

6 Conclusion

This paper develops a policy-friendly framework for analysing the interaction between aggregate demand and the dynamics of an epidemic. Our approach links a standard SEIR epidemiological model with models in the post-Keynesian and structuralist tradition, by allowing the infection rate to influence the utilisation rate and, in turn, the utilisation rate to affect the amount of physical contacts across individuals and hence the number of infections. We dub this the SEIR- u model.

We use the SEIR- u model to examine three different basic scenarios: laissez-faire, physical distancing without hard restrictions on economic activity, and lockdowns – modelled as government-mandated caps on the utilisation rate. A laissez-faire approach causes endogenous, sharp fluctuations of the predator-prey type in both utilisation and infections before the latter die out as herd immunity is reached in the long run. Under reasonable assumptions, physical distancing measures alone cannot reduce the reproduction rate under one, and the economy behaves in a manner qualitatively identical to the laissez-faire scenario. Lockdowns, in contrast, are very effective in reducing infections, albeit at the cost of sharp but temporary recessions. Even when the restrictions are lifted, and economic activity resumes, the infection rate remains very low for a rather long period.

The SEIR- u model sheds light on some important (and widely debated) issues. For, first, it calls into question the “lives versus livelihood” policy trade-off that has been emphasised in political circles as well as the popular press. Lockdowns clearly dominate all alternative courses of action in terms of reducing pandemic-induced deaths: hard caps on economic activity pull the break on the spread of the virus, thus avoiding the health system being overrun, and allow public authorities to buy time. Although this does entail a significant fall in economic activity, the loss in aggregate output caused by lockdown measures may be much more limited than commonly believed, given the independent negative effect of the epidemic on economic activity. The SEIR- u model suggests that the main effect of a lockdown is to make the downturn only slightly more pronounced, but significantly shorter compared with the alternative scenarios.

Second, a reduction in aggregate output is not the only cost of an epidemic, and of the measures to contain it; equally important are its distributive effects. Thus, we have generalised the SEIR- u model in order to examine the effect of the virus on aggregate income shares. Assuming that workers’ bargaining power, and thus the wage share, depend on economic activity, we have highlighted the importance of labour market policies in order to protect wage incomes. Indeed, employment protection, furlough programmes, and extensive provision of unemployment benefits are crucial especially if stringent public health measures affecting economic activity are adopted.

Of course, our analysis does not cover the full range of policy issues related to a major epidemic

– such as the SARS-CoV-2 pandemic – which has ramifications in virtually all aspects of social and economic life. Our main aim, however, is to provide a general framework to understand the interaction between epidemiological and economic variables within which such issues can be analysed. The SEIR- u model can be extended in a number of directions to tackle a range of questions. It outlines, we believe, a fruitful research programme.

References

- Atkeson, A., 2020. What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios. Working Paper 26867, National Bureau of Economic Research.
- Atkeson, A., Kopecky, K., Zha, T., 2020. Estimating and forecasting disease scenarios for COVID-19 with an SIR model. Working Paper 27335, National Bureau of Economic Research.
- Basurto, A., Dawid, H., Harting, P., Hepp, J., Kohlweyer, D., 2020. Economic and epidemic implications of virus containment policies: Insights from agent-based simulations. Working Paper 05-2020, Bielefeld.
- Bellomo, N., Bingham, R., Chaplain, M. A., Dosi, G., Forni, G., Knopoff, D. A., Lowengrub, J., Twarock, R., Virgillito, M. E., 2020. A multi-scale model of virus pandemic: Heterogeneous interactive entities in a globally connected world. LEM Papers Series 2020/16, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa.
- Caiani, A., Fagiolo, G., Ferraresi, T., Ghezzi, L., Guerini, M., Lamperti, F., Napoletano, M., Roven-
tini, A., Vanni, F., 2021. On the economic and health impact of the COVID-19 shock on Italian regions: A value-chain approach. LEM Papers Series forthcoming, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa.
- Chen, F., 2012. A mathematical analysis of public avoidance behavior during epidemics using game theory. *Journal of Theoretical Biology* 302, 18 – 28.
- Chen, F., Toxvaerd, F., 2014. The economics of vaccination. *Journal of Theoretical Biology* 363, 105 – 117.
- Chiarella, C., Flaschel, P., 2000. *The Dynamics of Keynesian Monetary Growth: Macrofoundations*. Cambridge University Press, Cambridge.
- Chiarella, C., Flaschel, P., Franke, R., 2005. *Foundations of a Disequilibrium Theory of the Business Cycle*. Cambridge University Press, Cambridge.
- Di Guilmi, C., Galanis, G., Baskozos, G., June 2020. A behavioural SIR model and its implications for physical distancing. Tech. rep., CRETA Online Discussion Paper Series 58.
- Eichenbaum, M. S., Rebelo, S., Trabandt, M., 2020. The macroeconomics of epidemics. Working Paper 26882, National Bureau of Economic Research.

- Eksin, C., Paarporn, K., Weitz, J. S., 2019. Systematic biases in disease forecasting. The role of behavior change. *Epidemics* 27, 96 – 105.
- Flaschel, P., Charpe, M., Galanis, G., Proano, C. R., Veneziani, R., 2018. Macroeconomic and stock market interactions with endogenous aggregate sentiment dynamics. *Journal of Economic Dynamics and Control* 91, 237 – 256.
- Flaxman, S., Mishra, S., Gandy, A., et al., 2020. Estimating the effects of non-pharmaceutical interventions on covid-19 in europe. *Nature* 584, 257 – 261.
- Galanis, G., Di Guilmi, C., Bennett, D. L., Baskozos, G., 2020. The effectiveness of non pharmaceutical interventions in reducing the outcomes of the COVID-19 epidemic in the UK, an observational and modelling study. Tech. rep., medRxiv.
- Gallo, E., Barbieri-Goes, M. C., 2020. Infection IS the cycle: Unemployment, output and economic policies in the COVID-19 pandemic. *Review of Political Economy* forthcoming.
- Goenka, A., Liu, L., Nguyen, M.-H., 2014. Infectious diseases and economic growth. *Journal of Mathematical Economics* 50, 34 – 53.
- Hale, T., Atav, T., Hallas, L., Phillips, T., Kira, B., Petherick, A., Pott, A., 2020a. Variation in US States' responses to COVID-19. Working Paper BSG-WP-2020/034, Blavatnik School of Government, University of Oxford.
- Hale, T., Webster, S., Petherick, A., Phillips, T., Kira, B., 2020b. Oxford COVID-19 Government Response Tracker. Tech. rep., Blavatnik School of Government.
- Hethcote, H. W., Dec. 2000. The mathematics of infectious diseases. *SIAM Rev.* 42 (4), 599 – 653.
- Kaplan, G., Moll, B., Violante, G. L., September 2020. The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the U.S. Working Paper 27794, National Bureau of Economic Research.
- Kermack, W. O., McKendrick, A. G., 1927. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London Series A* 115 (772), 700–721.
- Li, R., Pei, S., Chen, B., Song, Y., Zhang, T., Yang, W., Shaman, J., 2020. Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2) 368 (6490), 489–493.
- Manfredi, P., D'Onofrio, A., 2013. *Modelling the interplay between human behavior and the spread of infectious diseases*. Springer: New York.
- McFadden, D., 2001. Economic Choices. *American Economic Review* 91 (3), 351–378.

- Mitze, T., Kosfeld, R., Rode, J., Wälde, K., 2020. Face masks considerably reduce COVID-19 cases in Germany: A synthetic control method approach. Working Paper 13319, Institute of Labor Economics (IZA).
- Petherick, A., Beatriz, K., Lorena, B., Rodrigo, F., Maria, L., Saptarshi, M., Rafael, G., 2020. Brazil’s fight against COVID-19: risk, policies, and behaviours. Working Paper BSG-WP-2020/036, Blavatnik School of Government, University of Oxford.
- Piketty, T., 2014. Capital in the Twenty-First Century. Harvard University Press: Cambridge, MA.
- Prem, K., Flasche, S., et al., 2020. The effect of control strategies to reduce social mixing on outcomes of the covid-19 epidemic in wuhan, china: a modelling study. The Lancet Public Health 5, 261 – 270.
- Razmi, A., 2020. Some short-run macroeconomic considerations as society deals with a once-in-generations pandemic. Working Paper 2020-04, University of Massachusetts Amherst.
- Rowthorn, R., Toxvaerd, F. A., 2020. The optimal control of infectious diseases via prevention and treatment. Tech. Rep. 2020/13, University of Cambridge.
- Sethi, S., 1978. Optimal quarantine programmes for controlling an epidemic spread. Journal of the Operational Research Society 29, 265 – 268.
- Toxvaerd, F., 2019. Rational disinhibition and externalities in prevention. International Economic Review 60 (4), 1737–1755.
- Toxvaerd, F., 2020. Equilibrium social distancing. Working Paper CWPE2021, University of Cambridge.
- Train, K., 2009. Discrete choice methods with simulation. Cambridge University Press: Cambridge.

A Behavioural microfoundations

The SEIR- u model focuses on the structural aspects of the interaction between the epidemic and economic activity and does not explicitly deal with individual decision-making. This does not mean that the key relations are inconsistent with reasonable assumptions on individual behaviour. In this section, we illustrate this point by providing behavioural microfoundations to equation (8).

Assume a large population of heterogeneous agents who, at each point in time, choose between physical distancing (d) and not (nd). Their choice depends on the utility of the two alternatives: agent i chooses d if and only if

$$U_i = u_i^d - u_i^{nd} \geq 0, \quad (20)$$

where u_i^d, u_i^{nd} represent the utility of individual i when choosing, respectively, d or nd , and U_i is the difference between the two. Following [Di Guilmi et al. \(2020\)](#) who build on discrete choice models along the lines of [McFadden \(2001\)](#) and [Train \(2009\)](#), we assume that

$$U_i = V + \epsilon_i, \quad (21)$$

where V is the same across agents and depends on observable factors and ϵ_i varies across individuals. For a given set of measures, the relevant determinant of the common component of U_i in our model, V , is the infection rate, $I(t)$, which determines the probability of being infected.³³ In representative agent models, everyone makes the same choice; here, for a given $I(t) < N$, in general some agents will choose d , other nd .

Let $P^d(I(t))$ be the probability of a random individual choosing d given $I(t)$, where the functional form of $P^d(I(t))$ depends on the distribution of ϵ_i . Drawing on the empirical analysis of [Galanis et al. \(2020\)](#), for the UK, we postulate a linear relationship such that

$$P^d = \zeta I(t). \quad (22)$$

Physical distancing practices have a direct effect on aggregate demand Y^D such that

$$Y^D = a_y(1 - P^d)Y = a_y(1 - \zeta I(t))Y, \quad (23)$$

and

$$\dot{Y} = \chi(Y^D - Y) = \chi[a_y(1 - \zeta I(t))Y - Y] \quad (24)$$

or

$$\hat{Y} = \chi(a_y - 1) - \chi\zeta I(t). \quad (25)$$

For $\chi(a_y - 1) = n$ and $\chi\zeta = \alpha$, we get equation [\(8\)](#).

B Numerical calibration

This section describes the numerical calibration of our model used in the baseline simulations.

We have chosen the value for the basic reproduction number in a laissez-faire scenario with economic activity running at the normal level – $r(0) = \beta_0 \bar{u} / \gamma$ – to be equal to 3 as this value sits in between the values used in two key studies on the topic ([Flaxman et al., 2020](#); [Li et al., 2020](#)). Following [Atkeson et al. \(2020\)](#) and [Prem et al. \(2020\)](#) we set $1/\gamma = 7$, corresponding to an average duration of illness of 7 days. Assuming a normal utilisation rate $\bar{u} = 1/3$ – the long-run average for the output-capital ratio in advanced countries such as the United States ([Piketty, 2014](#)) – it follows that $\beta_0 = 1.27273$. This implies that $r(t) \leq 1$ when $u(t) \leq 0.112245$.³⁴

³³Obviously during a hard lockdown individual choices are severely restricted as the public is compelled to take distancing measures.

³⁴A two-third reduction in economic activity may seem rather dramatic. Yet, it is worth stressing that this is the value of the utilisation rate required to contain the infection *absent any policy measures and behavioural changes*.

Given the obvious difficulty in choosing specific values of the parameters capturing the speed of convergence of the economic dynamics, δ , and the effect of infections on economic activity, α , we have opted for some conservative estimates for the baseline simulations and checked the robustness of our results for a wide range of values. To be specific, we have set $\delta = .2$, which implies that the utilisation rate displays a reasonably fast self-correcting tendency around its normal value, and $\alpha = .025$, implying that a one percentage-point increase in the infection rate reduces the rate of change of the utilisation rate by a quarter of a percentage point all else equal. In other words, consistent with our conservative approach, in the baseline simulations we have assumed that, absent any restrictive policies, the pandemic does not affect economic activity too severely.

Finally, we have initialised all simulations using the same values of the state variables. In order to capture a situation corresponding to the very beginning of an epidemic, – before infections have any perceivable effect on the economy, – all of our simulations have been initialised assuming that the economy is at the normal level, $u(0) = \bar{u} = 1/3$, and the virus has infected a very small proportion of the population, $E(0) = .35 * 10^{-4}$, while there are no infectious or recovered individuals yet and the number of susceptible agents is determined as a residual ($S(0) = 100 - E(0)$).

C Effective reproduction rates

Recall that $r(t) \equiv \beta(t)/\gamma$ is the so-called *simple* reproduction rate, and let $r^e(t) = r(t)S(t)/N$ be the *effective* reproduction rate. An important feature of our framework is that, unlike the baseline SEIR model, both rates are endogenous, because they depend on the level of economic activity $u(t)$. In this section, we describe the dynamics of $r(t)$ and $r^e(t)$ in the main simulations.

The left panel of Figure 16 plots the simple reproduction rate corresponding to three different scenarios: low physical distancing, lockdown policies, and the interaction between lockdowns and low distancing.³⁵ The right panel displays the effective reproduction rate in *laissez-faire* as compared to both the low physical distancing and the lockdown policy scenario.

³⁵The simple reproduction rate in *laissez-faire* is omitted from the plot for the following reason. The model is initialised at $u(0) = \bar{u}$ with an $r(0) = 3$. Since the utilisation rate converges to \bar{u} in the SEIR model after the herd immunity threshold is crossed – given that both γ and β_0 are parameters – the simple reproduction rate in *laissez-faire* returns to 3 when the pandemic is over.

Figure 16: SEIR-u model: reproduction rates.

