Road to recovery: From extinguishing an epidemic to managing it *

Simon Loertscher[†] Ellen V. Muir[‡]

April 23, 2020

Abstract

Without widespread immunization of the population or complete eradication of the virus, the road to recovery from the current COVID-19 lockdowns will follow a path that strikes the difficult balance between the social and economic benefits of liberty and the death toll from the disease. We provide an approach that combines epidemiology and economic models by taking as given the constraint that the maximum capacity of the health care system must not be exceeded. Treating the transmission rate as a decreasing function of the severity of the lockdown, we determine the minimal lockdown that satisfies this constraint.

Keywords: COVID-19, SIR models, capacity constraints, managing an epidemic

JEL-Classification: H51, I18

*We are grateful for suggestions and comments by Gabriel Carroll, Chris Edmond, Ian Harper, Zi Yang Kang, Paul Milgrom, Gary Stoneham and the theory seminar audience at the University of Melbourne. Financial support by the June and Samuel Hordern Endowment is also gratefully acknowledged.

[†]Department of Economics & Centre for Market Design, Level 4, FBE Building, 111 Barry Street, University of Melbourne, Victoria 3010, Australia. Email: simonl@unimelb.edu.au.

[‡]Department of Economics, Stanford University. Email: evmuir@stanford.edu

1 Introduction

Without widespread immunization of the population (for example, due to the development of a vaccine) or complete eradication of the virus, the road to recovery from the ongoing coronavirus-induced lockdowns will require sustained vigilance to ensure that the spread of the virus remains at a manageable level for a country's or region's healthcare system. At the same time, this recovery ought to start as soon as possible for a number of reasons, which include the curtailing of liberty that these lockdowns impose, the mental and other health issues associated with social distancing and isolation, and the tremendous economic cost that has no parallel in living memory. In other words, recovery requires the transition from a paradigm in which eradication of an epidemic is the goal to one in which the epidemic is *managed*.

In this article, we show how this can be done by providing a methodology that permits return to some kind of normalcy while keeping the spread of the virus at a level that even at the peak of the endemic does not exceed the capacity constraint of the health care system. Specifically, we use a standard epidemiology model—a simple *SIR (Susceptible-Infectious-Recovered) model*—to predict the peak of the epidemic while treating the rate of transmission as the variable that the policymaker can influence by choosing the severity of a lockdown. We treat as a hard constraint the capacity of the health care system, that is, the maximum number of COVID-19 patients that it can handle per period of time at the peak of the crisis.[1] Of course, this capacity constraint will need to be defined in such a way that patients with other—but no less severe—needs for care are still able to access treatment.[2]

The main contribution of this paper is to formulate an operational constraint that provides policymakers with guidance for how to manage an epidemic which is too costly to eradicate; to incorporate this constraint into a standard epidemiology model; and derive implications for the severity of the lockdown that is necessary to respect that constraint. It serves the purpose of a proof of concept. The model can and will need to be refined

[1]To fix ideas, throughout the paper we will talk about the capacity of the healthcare system as our binding constraint. However, the framework outlined in this paper can accommodate any constraint that can be expressed as a function of the number of COVID-19 cases that occur at the peak of the epidemic. This constraint can be interpreted as a normative criterion that society has to choose. For example, suppose that society viewed the number of deaths that would occur as a result of using the capacity of the healthcare system as a binding constraint as unacceptable. Then one could instead treat a fixed proportion of the healthcare system that is utilized at the height of the crisis as a binding constraint (such as requiring that the healthcare system never exceeds 80% capacity). An alternative, but mathematically equivalent, approach would be to place an upper bound on the number of deaths per day at the peak of the crisis, which may be in line with some of the concurrent debates (see, e.g., the New York Times article "The Cold Calculations America's Leaders Will Have to Make Before Reopening").

[2]In a public health catastrophe, this is not always the case; see, for example, this New York Times article: "The Pandemic's Hidden Victims: Sick or Dying, but Not From the Virus'.'

subsequently.

The data required to apply the proposed approach are accurate estimates of the transmission rate, reliable measures of the spread of the disease, and the percentage of the infected persons that require treatment. Some of these data can and will need to be obtained using ongoing random sampling of the population. Importantly, for practical purposes what we propose can be implemented using continuous, real-time updating and adjustments. To the extent that the population is continuously and randomly sampled and tested, decision makers will have a reasonably clear real-time picture of the spread of the virus, which help mitigate the problems of time-lag and inertia.

To convey a sense of the magnitude of the potential economic and social costs, consider the unemployment rate during the Great Depression in the U.S., then and now the world's largest economy, and the weekly unemployment filings in the U.S. in the wake of the ongoing coronavirus-related lockdown in 2020.[3] The immediate consequences of the Great Depression were mass poverty and economic devastation, and at least indirectly, the rise of Hitler in Germany.[4] The unemployment rate in the U.S. was 4.4% in March 2020 and is by many analysts expected to rise to 17% or above in April 2020. This means that an increase in the unemployment rate of roughly 12.5 percentage points—which during the Great Depression took two years to materialize—is expected to happen within a month.

Great D.	1929	1930	1931	1932	1933
Unemployment rate	3.2	8.7	15.9	23.6	24.9

2020	Mar 14	Mar 21	Mar 28	Apr 4	Apr 11
Unemployment filings	282	3,307	6,867	6,615	5,245

Not surprisingly, there has been a recent upsurge of interest in SIR models in economics. Atkeson (2020) provides an introduction of this modeling approach to economics, which is standard in mathematical biology (see, e.g., Murray, 2002). Alvarez et al. (2020) apply an optimal control approach to an SIR model to derive the optimal lockdown policy that trades off the cost of death against economic output. Our approach is similar to theirs in that we derive an optimal lockdown policy. In contrast to Alvarez et al. (2020), in our approach this policy is the minimal lockdown necessary to satisfy the constraint that the capacity of the health care system not be exceeded. As such, it takes as given that capacity constraint and does not involve an explicit tradeoff between lives saved and output. In recent

[3]Sources: thebalance.com/Reinhart and Rogoff (2009) and edition.cnn.com/US Department of Labor. Prior to March 21, 2020, the record number of weekly unemployment filings was 695,000 in October 1982 (i.e. shortly before the peak of the early 1980s recession).

[4]See Appendix A for a short discussion of the Great Depression and comparisons with today's economic data.

days and weeks, there has also been a rise in more informal commentary and discussions of the problems at hand (see, e.g., Gilbert et al., 2020) and analyses of tradeoffs involving economics without explicitly embedding epidemiology models such as Hall et al. (2020) and Budish (2020).

The remainder of this paper is organized as follows. Section 2 describes the setup. Section 3 derives the dynamics of an epidemic and the lockdown necessary to keep it at a level that respects the capacity constraint. Section 4 provides a discussion of possible and natural extensions of the baseline model we study, which serves as a proof of concept. Section 5 concludes the paper.

2 Setup

Consider a basic susceptible-infectious-recovered (SIR) model with a constant population of size N. This is a classic model in epidemiology (see, e.g., Murray, 2002, p.320), in which the population is divided into three compartments consisting of susceptible individuals, infected individuals and recovered individuals, respectively denoted by $S(t)$, $I(t)$ and $R(t)$ at time t. Note that because of the assumption of a constant population of size N, for all $t \geq 0$, we have

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

We let $N_1 = S(0)$, $N_2 = I(0)$ and $N_3 = R(0)$ and assume that only two types of transitions are possible: susceptible individuals can become infected and infected individuals recover. (As is standard, "recovered" simply means the individuals are no longer infectious, which occurs either because they gained immunity or died following infection.) We let β denote the average number of contacts per person per time and assume that we have a well-mixed or homogeneous population so that $I(t)/N$ is the fraction of contact occurrences that involve an infectious individual. The *rate of transition* between the susceptible compartment and the infectious compartment is thus given by $\beta I(t)/N$.

We denote by $\ell \in [0, 1]$ the severity of the lockdown, with $\ell = 0$ meaning no lockdown and $\ell = 1$ meaning complete lockdown. We assume that ℓ is the choice variable of the policymaker, and with regards to the epidemic, its impact is that it affects the transmission rate β as follows:

$$\beta(\ell) = \beta_0 + (1 - \ell)\beta_1,$$

where $\beta_0 \geq 0$ is a fixed component of the transmission rate, $\beta_1 > 0$ is a constant, and $\beta(\ell)$ makes the dependence of β on ℓ explicit.

We further assume that individuals recover at rate γ.[5] In SIR models, the parameter $R_0 = \beta/\gamma$ plays an important role in governing the dynamics of an epidemic. In this simple version whenever $R_0 N_1 > 1$ the number of infected individuals will increase from time $t = 0$, resulting in an *epidemic*. If $R_0 N_1 < 1$ then the number of infected individuals will decrease from time $t = 0$ and an epidemic does not occur (alternatively, we can think of the "peak" of the epidemic as occurring at time $t = 0$).

The proportion $\tau \in [0, 1]$ of those who are infected need *treatment*, so that, given $I(t)$ and τ, the number of people requiring treatment at time t is

$$T(t) = \tau I(t).$$

Letting $K > 0$ denote the maximum capacity of the health care system to treat COVID-19 patients without reducing the care given to other patients in need, the constraint for managing the epidemic is, for all $t \geq 0$,

$$T(t) \leq K. \tag{1}$$

In Section 3.4, we augment the epidemiology model by an economic production function to analyze tradeoffs involving economics. Specifically, we assume that GDP, denoted Y, is produced using labor L according to the production function $Y = L^\alpha$, where $\alpha \in (0, 1)$ is a parameter that measures labor's productivity, which can be calibrated using labor's income share in national accounts data.[6] Letting $L_0 \geq 0$ denote the amount of labor that is not affected by the lockdown variable ℓ, the amount of labor that is productive given ℓ is

$$L(\ell) = L_0 + (1 - \ell)L_1,$$

where L_1 is the part of the labor that is affected by the lockdown variable ℓ. It follows that $L(0)$ is the pre-lockdown labor supply.

3 Managing an epidemic

We now analyze the dynamics of an epidemic and then derive the minimal lockdown policy ℓ_K necessary to satisfy the constraint K at all times.

[5]For example, if D is the duration of infection then the rate of recovery is given by $\gamma = 1/D$.

[6]This corresponds to assuming a Cobb-Douglas production function with all input factors other than labor being fixed for the duration of the disease; see, for example, Jehle and Reny (2011).

3.1 Dynamics of an epidemic

The dynamics of an epidemic in our simple SIR model are governed by the following system of non-linear differential equations:

$$\frac{dS(t)}{dt} = -\frac{\beta I(t)S(t)}{N}, \quad \frac{dI(t)}{dt} = \frac{\beta I(t)S(t)}{N} - \gamma I(t) \quad \text{and} \quad \frac{dR(t)}{dt} = \gamma I(t),$$

with initial conditions $S(0) = N_1$, $I(0) = N_2$ and $R(0) = N_3$. Harko et al. (2014) provided an analytic solution to this system of equations by parameterizing time t by a parameter u. In particular, introducing the integration constants

$$S_0 = N_1 \mathrm{e}^{\frac{\beta N_3}{\gamma}}, \quad u_0 = \mathrm{e}^{-\frac{\beta N_3}{\gamma}}, \quad \text{and} \quad C_1 = -\beta N$$

we have

$$t(u) = \int_{u_0}^{u} \frac{1}{\xi(C_1 - \gamma \log(\xi) + S_0 \beta \xi)} d\xi, \; I(u) = N - S_0 u + \frac{\gamma \log(u)}{\beta} \text{ and } R(u) = -\frac{\gamma \log(u)}{\beta}.$$

Notice that when $u = u_0$ we have $t = 0$ and that u decreases as t increases.[7]

The basic dynamics of an epidemic are as follows. As susceptible individuals become infected and then recover, the stock of susceptible individuals decreases over time and the stock of recovered individuals increases over time. The number of infected individuals initially increases before reaching an epidemic peak and then gradually decreasing. The number of infected individuals stops increasing once the population of susceptible individuals is sufficiently small. An example of a typical epidemic path is shown in Figure 1. Note that unless stated otherwise all figures are drawn for the parameterization $N_1 = 0.999$, $N_2 = 0.001$, $N_3 = 0$, and $\gamma = 1/18$; Figure 1 assumes $\beta = 3.1/18$.[8]

Assuming $R_0 N_1 = \beta N_1/\gamma > 1$ (so that the peak of the epidemic does not occur at $t = 0$), the maximal number of infected individuals $I^*(\beta)$ during the epidemic is characterized by

$$I^*(\beta) = N - S_0 u_{\max} + \frac{\gamma \log(u_{\max})}{\beta},$$

[7] In the limit as $t \to \infty$ we have $u \to -\gamma W\left(-\mathrm{e}^{\frac{C_1}{\gamma}} \beta S_0/\gamma\right)/(\beta S_0)$, where W is the product log function.

[8] Here, we normalize the size of the population to 1 and assume that initially 0.1% of the population is infected. Following Wang et al. (2020), we take $\gamma = 1/18$, which reflects an average disease duration of 18 days (and so the appropriate interpretation of the time scale t is then also in days). We also set $\beta = \beta_0$ so that $R_0 = 3.1$, which is in line with estimates from Wuhan, China prior to the introduction of strict lockdown measures.

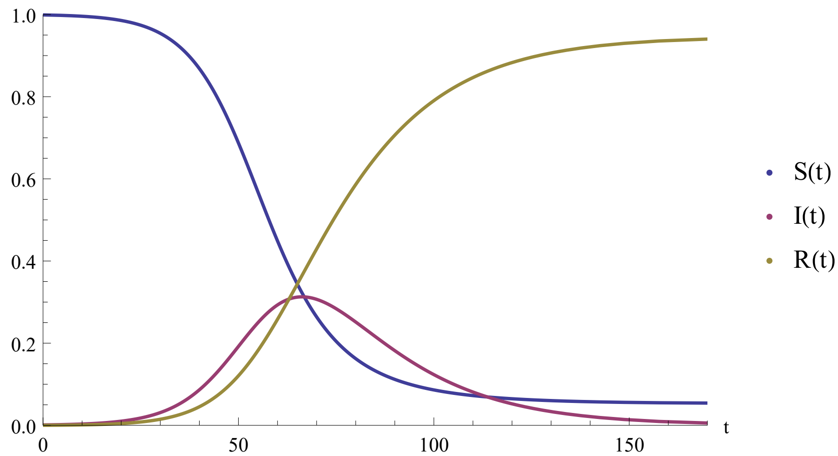

Figure 1: The evolution of a typical epidemic.

where

$$u_{\max} = \frac{\gamma}{\beta S_0}.$$

Notice that we have

$$\frac{d\,I^*(\beta)}{d\,\beta} = -\frac{\gamma \log\left(\frac{\gamma}{\beta N_1}\right)}{\beta^2} > 0, \tag{2}$$

where the inequality follows from the fact that $\log(\gamma/(\beta N_1)) < 0$ since by assumption $\gamma/(\beta N_1) < 1$.

3.2 Binding, slack or violated capacity constraints

Notice that the parameters β_0, β_1, γ and τ, as well as the initial conditions, impose restrictions on the lower feasible bound for K. Specifically, denote by $I_\ell(t)$ the number of infectious at time t given policy $\ell \in [0, 1]$, by

$$I^*_\ell = I^*(\beta(\ell))$$

the maximum number of infected individuals given ℓ, and by

$$T^*_\ell = \tau I^*_\ell$$

the maximum number of people needing treatment per time given policy ℓ. From (2) we have that I^*_ℓ and T^*_ℓ are continuously decreasing in ℓ since $\beta(\ell)$ is decreasing in ℓ. From this and continuity follows that K is feasible if and only if

$$K \in [T^*_1, T^*_0]. \tag{3}$$

If $K < T_1^*$, then the capacity constraint is so tight that it can never be satisfied at the peak of the endemic, not even with the most severe lockdown policy. If $K > T_0^*$, then no lockdown is required to satisfy the constraint.

Conversely, for any K satisfying (3) there is a *minimal lockdown policy*, denoted ℓ_K, that satisfies the constraint that the number of individuals requiring treatment at time t never exceeds K. Formally,

$$\ell_K := \min\{\ell \mid \ell \in [0,1] \quad \text{and} \quad K \leq T_\ell^*\}.$$

Because T_ℓ^* is a decreasing functions of ℓ, it follows that ℓ_K is a decreasing function of K. Intuitively, as the capacity constraint K increases, the severity of the required lockdown decreases.

As for policy implications, this means that, all else equal, states or countries with larger capacities can afford less stringent lockdowns. For a given lockdown policy the transmission rate parameter β can also vary substantively between states and countries as the value of this parameter varies with factors such as population density and household composition. Since the maximum number of patients requiring treatment is given by $\tau I^*(\beta)$ and $I^*(\beta)$ is increasing in β (see (2)), it follows that, all else equal, states or countries with larger transmission rates require more stringent lockdowns. Formally, compare two regions, each with capacity K, with transmission rates parameterized by (β_0, β_1) and $(\hat{\beta}_0, \hat{\beta}_1)$ satisfying $\hat{\beta}_i \geq \beta_i$ for $i = 0, 1$, where at least one of these inequalities is strict. Denoting the respective minimal lockdown policies by ℓ_K and $\hat{\ell}_K$, we then have

$$\ell_K < \hat{\ell}_K.$$

In other words, regions with lower transmission rates can afford slacker lockdown policies as is illustrated in Panel (a) of Figure 2. This figure uses the same parameters values as Figure 1 but with $(\beta_0, \beta_1) = (0.5/18, 2.6/18)$ and $(\hat{\beta}_0, \hat{\beta}_1) = (0.6/18, 2.6/18)$.

3.3 The relationship between lockdown and capacity

We now look in slightly more detail at the relationship between K and ℓ_K. Panel (b) of Figure 2 plots this relationship, assuming $\beta(\ell) = 0.5\gamma + 2.6\gamma(1-\ell)$ and $\tau = 0.11$.[9] As

[9] Following Wang et al. (2020), we use $R_0 = 3.1$ with no lockdown and $R_0 = 0.5$ under the strictest possible lockdown. The first of these R_0 values is an estimate for Wuhan, China prior to any policy interventions by the Chinese government. The second of these R_0 values is an estimate for Wuhan, China under the strictest lockdown measures implemented by the government. We use $\tau = 0.11$, which is consistent with data from New York which showed that around 11% of confirmed coronavirus patients were hospitalized at

(a) A range of γ values

(b) A range of τ values

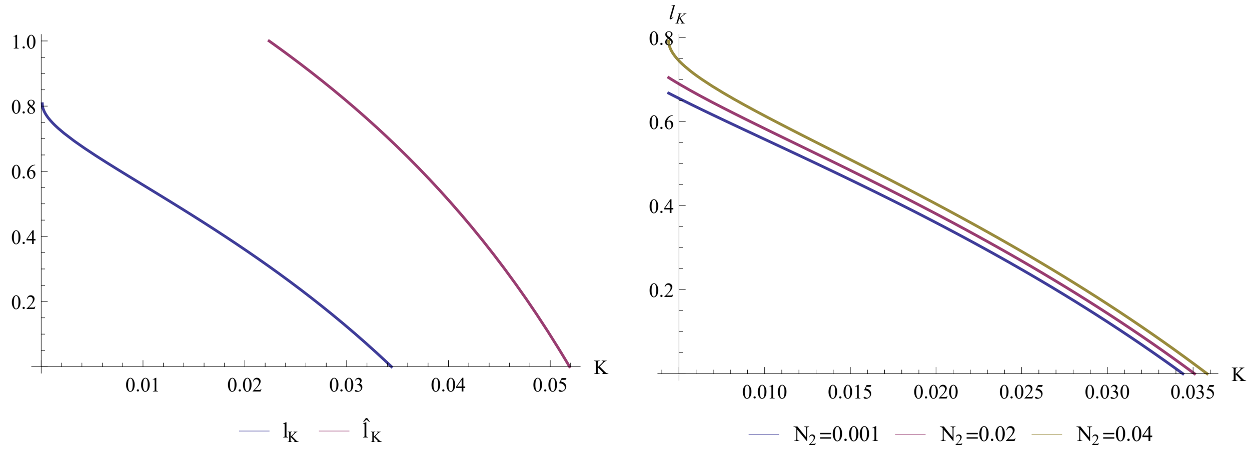

Figure 2: Panel (a) illustrates that a higher schedule of β values necessitates a more severe lockdown for a given K value. Panel (b) illustrates the relationship between ℓ_K and K for a range of N_2 values. As was shown analytically, for a given N_2 value, the severity of the lockdown ℓ_K decreases as the capacity K of the healthcare system increases. This figure also shows that a more severe lockdown is required if a higher proportion N_2 of the population is initially infected.

before we set $N_3 = 0$ and $\gamma = 1/18$ but we now create plots for three different values of N_2: $N_2 = 0.001$ (in which case $N_1 = 0.999$), $N_2 = 0.02$ (in which case $N_1 = 0.98$) and $N_1 = 0.04$ (in which case $N_1 = 0.96$). Panel (b) of Figure 2 shows that the lockdown policy ℓ needed to achieve a given K increases in the proportion of the population that is initially infected. This figure also shows how the proportion of the population that requires treatment at the height of the pandemic, for a given lockdown policy ℓ, increases in the proportion of individuals N_2 that is initially infected. Consequently, for a given cap K, a more severe lockdown is required as N_2 increases. This result highlights the high cost of a delayed policy response.[10]

Figure 3 provides some additional comparative statics showing how the severity of the

the peak in hospitalizations (Feuer, 2020). Note that τ is not the rate of hospitalization (i.e. the proportion of coronavirus patients that are hospitalized at some point over the course of their illness) but rather the proportion of coronavirus patients that are hospitalized at any given point in time. An alternative approach would be to include a separate compartment in the model for hospitalizations, the main advantage being that this would produce a time-lag between the peak in the number of infected individuals and the peak in the number of hospitalized individuals (which is consistent with what we observe in data). We ensure that the binding constraint on the healthcare system is not violated due to this time-lag effect by calibrating the model using the proportion of the population that is hospitalized at the peak in hospitalizations.

[10]For example, when San Francisco issued a shelter-in-place order on March 16, 2020 its number of per capita confirmed coronavirus cases was comparable to that of New York City. By the time New York City was subject to a shelter-in-place order on March 23, 2020 the number of per capita confirmed coronavirus cases was greater than that of San Francisco by roughly an order of magnitude.

(a) A range of γ values

(b) A range of τ values

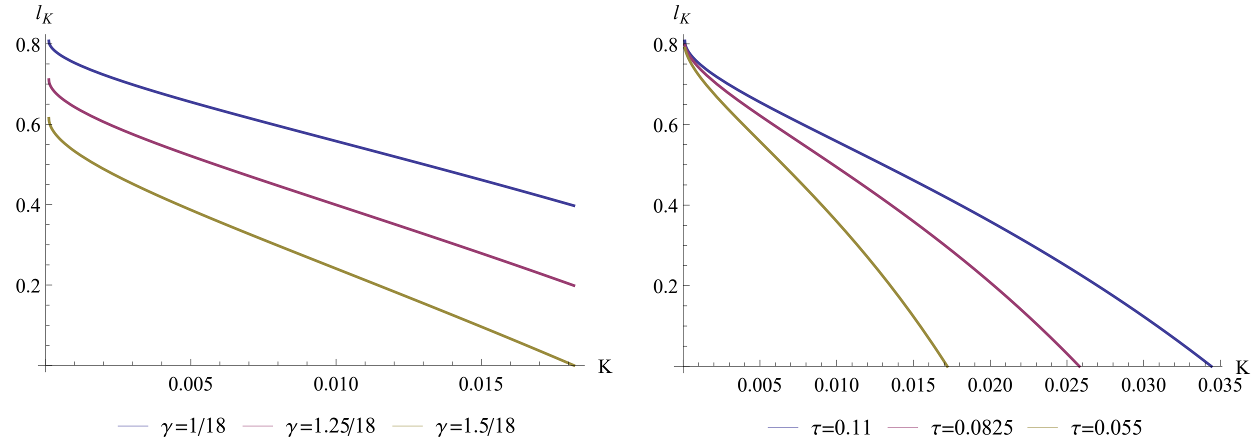

Figure 3: The severity of the required lockdown decreases as γ increases and as τ decreases.

required lockdown decreases as γ increases and as τ decreases. Panel (a) uses the same parameter values as Panel (b) of Figure 2 but with $N_2 = 0.001$ (in which case $N_2 = 0.999$ since we set $N_3 = 0$) and $\gamma = 1/18$, $1.25/18$ and $1.5/18$. Panel (b) also uses the same parameter values as Panel (b) of Figure 2 but with $N_2 = 0.001$ (in which case $N_2 = 0.999$ since we set $N_3 = 0$) and $\tau = 0.11$, 0.0875 and 0.055. One interpretation of these comparative statics is that as superior treatments become available, individuals both require less overall treatment and recover from the disease more quickly and hence a less severe lockdown is required.

3.4 Economic impact

By mapping the severity of the lockdown to gross domestic product (GDP), one can trace out the relationship between the capacity constraint and economic output. Substituting $L(\ell)$ into the production function yields output

$$Y(\ell) = (L_0 + (1-\ell)L_1)^{\alpha}.$$

It follows that, given the minimal lockdown policy ℓ_K for the constraint K, output, denoted Y_K, is given by

$$Y_K = Y(\ell_K).$$

Because $Y(\ell)$ decreases in ℓ and ℓ_K decreases in K, it follows that

$$\frac{dY_K}{dK} > 0.$$

A plot illustrating how Y_K increases in K for a given set of parameters can be found in Figure 4, which assumes $L_0 = 1/3$, $L_1 = 2/3$ and $\alpha = 1/3$ (and as before $N_1 = 0.001, 0.02, 0.04$, $N_3 = 0$, $\beta(\ell) = 0.5\gamma + 2.6\gamma(1-\ell)$, $\gamma = 1/18$, $\tau = 0.11$).[11] This plot shows that longer delay in the initial policy response—which leads to a higher number of infected individuals in the population prior to any lockdown intervention—results in policy makers facing a more severe economic impact of the pandemic in order to satisfy the binding constraint K.

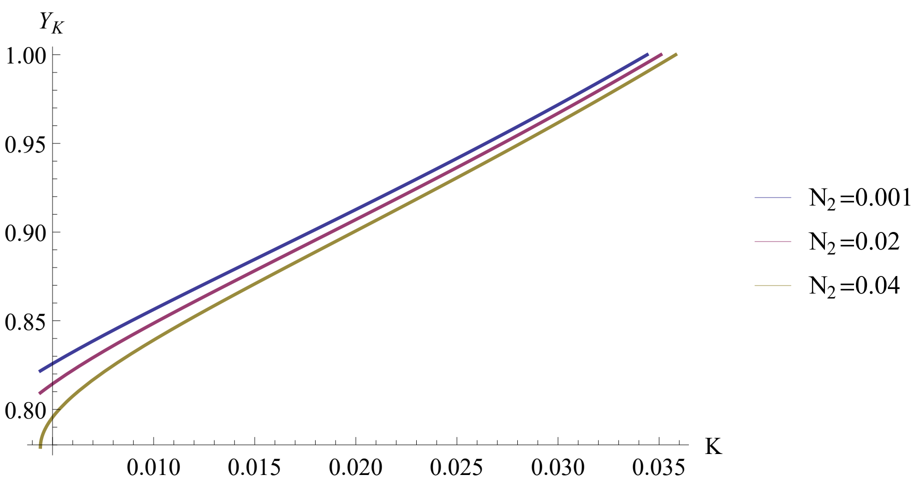

Figure 4: Output Y_K increases in K and decreases in N_2.

4 Discussion

We now provide a brief discussion. We first elaborate on our epidemiology modeling approach and natural alternatives and extensions. Then we address sensitivity of the model and the need and possibility to construct confidence intervals for the peak of the epidemic. The section concludes with a short discussion of dynamically adjusting policies.

[11] We use a standard value of $\alpha = 1/3$ and setting $L_0 = 1/3$ corresponds to a labor force in which one third of workers are essential. While we do not pursue this here, our model also lends itself to the possibility of relating the costs and benefits of a lockdown to the statistical value of life such as DALY (disability-adjusted life year) or QALY (quality-adjusted life year) that are widely used in public health debates. See Hall et al. (2020) for a framework that considers the tradeoff between consumption and COVID-19 deaths. By expressing the value of a life in terms of years of per capita consumption this approach allows those authors to derive an upper bound on the level of consumption a utilitarian society would be willing to forgo in order to avoid COVID-19 deaths.

4.1 Epidemiology modelling alternatives

The model we have analyzed is one with a constant population. A first natural extension would be to include births and deaths, which are sometimes also referred to as *vital dynamics*. Moreover, one could include additional compartments—such as having a compartment for exposed individuals and another for hospitalized individuals—or additional transitions by allowing, for example, for the possibility that recovered individuals become susceptible.

However, regarding the COVID-19 crisis, there is reason to hope that a vaccine will become available sometime in 2021. Consequently, the main concern among policy makers is appropriately responding to the first wave of infections. As long as recovered individuals retain immunity until the introduction of a vaccine, it would therefore not be necessary to account for the possibility of recovered individuals becoming susceptible again in models that are designed to deal with the current crisis.[12] Similarly, assuming a constant population is probably a reasonable approximation for the problem at hand, where a relatively short time period is relevant and the death rate for individuals that participate in the workforce is very low. Accounting for the impact of births (including immigration), deaths (including those from coronavirus) and aging (i.e. people moving into and out of the labor force) on the dynamics of the epidemic as well as the size of the labor force would greatly complicate the model without substantively changing its predictions.

That said, an important and useful extension would be to account for the age structure of the population (particularly how susceptibility and hospitalization vary with age) as well as that of the labor force. This would allow us to evaluate policies such as allowing low-risk age groups (such as 20- to 30-year-olds who do not reside in households with vulnerable individuals) to return to work.[13] Such policies could help thwart or mitigate an economic catastrophe without causing a large increase in the spread of the virus. The ability to simultaneously quantify the impact of such policies on both the healthcare system and the economy illustrates the importance and usefulness of combining epidemiology and economic models.

Continuum SIR models, such as the one analyzed here, provide good approximations for large populations. However, for smaller populations or more refined targets—like ensuring that ICU beds do not run out or targeting a relatively narrow death rate band—this family of models does not necessarily provide good approximations. For these kinds of applications,

[12]Accounting for individuals transition from recovered back to susceptible is important, for example, when studying the seasonal dynamics of diseases such as influenza.

[13]Verity et al. (2020, Table 1 (last column)) report dramatic differences in age-specific fatality-to-infected ratios. For example, for the group of individuals 60 years old and older, this ratio is 3.28% while for the cohort of individuals in their 20s, it is 0.0309%, that is, the individuals in their 20s are roughly 100 times less likely to die when infected than those past 60.

models using agent-based simulations are more appropriate tools. In simulation-based approaches, it is relatively straightforward to include extensions such as age-structuring the population and labor force because there is no need to seek analytic solutions.

4.2 Sensitivity analysis and confidence intervals

Up to this point we have treated the parameters of the model as given and known. In practice, there may be considerable uncertainty and measurement error associated with these parameter values, implying that the predictions of the model are not deterministic. We now discuss how the distribution of the predicted peak of the epidemic can be derived from the, by assumption, known distributions of the uncertain parameters R_0, τ, N_1, N_2 and N_3.

After some tedious algebra and imposing the normalization $N = 1$, the proportion of the population requiring treatment at the peak of the epidemic is given by

$$T^*(R_0, \tau, N_1, N_3) = \tau \left(1 - N_3 - \frac{\log(R_0) + \log(N_1) + 1}{R_0}\right).$$

The density of T^* is thus

$$f_{T^*}(y) = \int_{\mathbb{R}^5} f_{R_0}(x_1) f_\tau(x_2) f_{N_1}(x_3) f_{N_3}(x_4) \delta(y - T^*(x_1, x_2, x_3, x_4))\, dx_1\, dx_2\, dx_3\, dx_4,$$

where δ denotes the Dirac delta function and f_X denotes the distribution of the random variable X. From here one can construct a confidence interval for the maximum number of individuals requiring treatment at the peak of the epidemic.

For example, suppose that R_0, N_1 and N_3 are known parameters and that $\tau \sim \mathrm{N}(0.11, 0.01)$. That is, τ is normally distributed with mean 0.11 and standard deviation 0.01. Then $T^*(R_0, \tau, N_1, N_3)$ is normally distributed with mean $\mu(R_0, N_1, N_2) = T^*(R_0, 0.11, N_1, N_3)$ and standard deviation

$$\sigma(R_0, N_1, N_2) = 0.01 \left(1 - N_3 - \frac{\log(R_0) + \log(N_1) + 1}{R_0}\right).$$

A 95% confidence interval for the value of T^* is then given by $[\mu - 1.96\sigma, \mu + 1.96\sigma]$. Therefore, if we use $T = \mu(R_0, N_1, N_2) + 1.96\sigma(R_0, N_1, N_2)$ then we can say that with 97.5% confidence, the constraint K will not be violated at the peak of the epidemic. An illustration is provided in Figure 5. This figure uses precisely the same parameters as those shown in Panel (b) of Figure 2 and sets $N_2 = 0.001$ (in which case $N_1 = 0.999$).

An alternative to the approach adopted here would be to perform a conditional worst-case

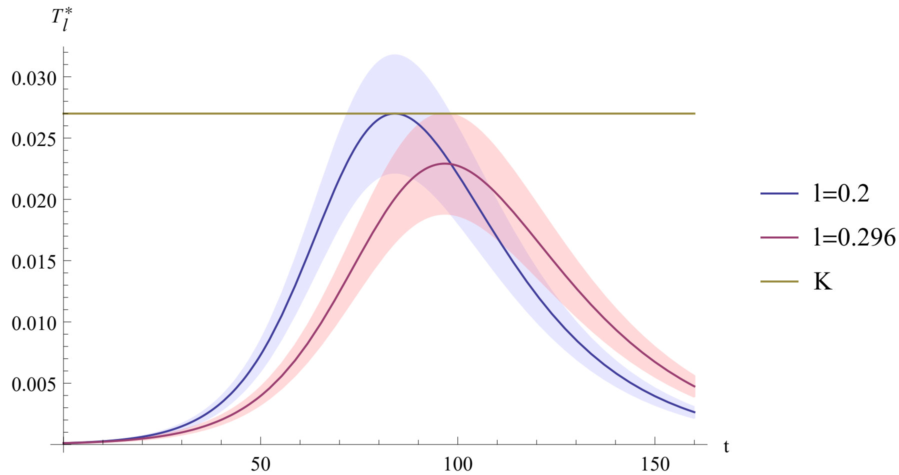

Figure 5: Assume $K = 0.0270$. If τ is deterministic and equal to 0.11, we have $\ell_K = 0.2$. In contrast, if τ is normally distributed with mean 0.11 and a standard deviation of 0.01, then $\ell = 0.296$ is required to satisfy the constraint with a probability of 0.975.

analysis. That is, if one had estimates of moments of a particular parameter distribution (such as its mean and variance), then one could compute confidence intervals with respect to the worst-case distribution.

4.3 Dynamically adjusting policies

Another natural extension of the theoretical analysis would be to allow the lockdown policy to vary over time and solve the control problem that pins down the optimal dynamic lockdown policy. For example, following the peak of the epidemic a less stringent lockdown would suffice in order to ensure that the capacity constraint remains satisfied. This analysis would also serve to shed light on issues such as whether a shorter and more severe lockdown would allow for a higher overall level of economic output without violating the capacity of the healthcare system, relative to a less severe but more protracted lockdown.

Similarly, in practice, new and arguably more accurate parameter estimates will probably be obtained as the epidemic progresses, which will allow policy makers to adjust policies dynamically.

5 Conclusions

This time *is* different.[14] The cause of the economic downturn (a pandemic rather than the burst of a financial bubble or any other structural issue with the economy), its scope (universal, hitting all countries more or less within the same quarter) and magnitude (record increases in unemployment filings in the United States) are unprecedented. While there are good reasons to be confident that, informed by the in-depth analyses of past mistakes, the policy response to a severe economic downturn will be better and swifter than at the onset and during the Great Depression, the unparalleled nature of the current shock makes recovery a perilous and winding road. Although policy makers are ready to act swiftly, the ongoing virulence of the disease may prevent them from so doing. Without widespread immunization, return to normalcy would be difficult if not impossible even if there were no inertia in rebooting economies that have come to a standstill. We will have to find the path to recovery by learning on the go, and learning quickly.

This paper proposes a pathway to recovery. It suggests a shift in thinking from extinguishing an epidemic in its entirety at all costs, to instead thoughtfully managing it with regards to the tradeoffs between managing the pandemic and economic output. This can be achieved by targeting an appropriately chosen constraint, such as the capacity of the healthcare system. Managing such a non-stationary process is unusual for economists while aiming for non-zero numbers of infections is also new territory for epidemiologists. It will be akin to learning to balance on a rolling ball.

The proposed pathway requires countries and other relevant jurisdictions to get ready for randomly testing the population in ways that are representative of this population in order to obtain an accurate picture of the spread of the disease in real time. (This will require a shift in attitude by public health officials and personnel from addressing the immediate needs of a public health crisis toward managing a wider economic crisis that has a public health component and cause. Rather than using testing kits exclusively on individuals suspected of carrying the virus, it will require deliberately using a proportion of these scarce kits on people from the general population for whom there is no particular reason to believe that they are infected.) It also necessitates the development of monitoring approaches that permit for intermediate levels of lockdowns. Although these specific challenges are not exactly common, they seem practical enough to be well within the grasp of any diligent government.

[14]This sentence is the title of the New York Times bestseller by Reinhart and Rogoff (2009), where it is used to explain why financial crises occur—because decision makers, in the lead-up to a financial crisis, trend to ignore important precedents. Here and now, however, it seems an accurate description of the current COVID-19 crisis.

References

Alvarez, F. E., D. Argente, and F. Lippi (2020): "A Simple Planning Problem for Covid-19 Lockdown," NBER Working Paper # 26981.

Atkeson, A. (2020): "What Will Be the Economic Impact of Covid-19 in the US? Rough Estimates of Disease Scenarios," NBER Working Paper # 26867.

Budish, E. (2020): "$R < 1$ as an Economic Constraint: Can We "Expand the Frontier" in the Fight Against Covid-19?" Note.

Feuer, A. (2020): "Stark Death Toll, but Cautious Optimism in N.Y. Over Hospitalizations," *New York Times*.

Gilbert, M., M. Dewatripont, E. Muraille, J.-P. Platteau, and M. Goldman (2020): "Preparing for a responsible lockdown exit strategy," *Nature Medicine*, Published Online, April 14.

Hall, R. E., C. I. Jones, and P. J. Klenow (2020): "Trading Off Consumption and COVID-19 Deaths," Working paper.

Harko, T., F. S. N. Lobo, and M. K. Mak (2014): "Exact analytical solutions of the Susceptible-Infected-Recovered (SIR) epidemic model and of the SIR model with equal death and birth rates," *Applied Mathematics and Computation*, 236, 184–194.

Jehle, G. and P. Reny (2011): *Advanced Microeconomic Theory*, Addison-Wesley, 3rd ed.

Murray, J. D. (2002): *Mathematical Biology: I. An Introduction*, Berlin: Springer, 3rd ed.

Reinhart, C. and K. Rogoff (2009): *This time is different*, Princeton, New Jersey: Princeton University Press.

Verity, R., L. C. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, ..., and N. M. Ferguson (2020): "Estimates of the severity of coronavirus disease 2019: a model-based analysis," *The Lancet*, Published Online, March 30.

Wang, H., Z. Wang, Y. Dong, R. Chang, C. Xu, X. Yu, ..., and Y. Cai (2020): "Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China," *Cell Discovery*, 6.

Appendix

A The Great Depression

Economists typically are seriously concerned when an ongoing crisis is compared to the Great Depression. The immediate consequences of the Great Depression were mass poverty and economic devastation. There are ample descriptions in popular literature and culture of the impact of this downturn on the lives of millions of people. At least indirectly, it arguably led to the rise of Hitler in Germany and thereby to WWII. A look at key data is insightful for comparisons. Table 1 contains unemployment rates and GDP growth for U.S. for the years 1930 to 1933. Table 2 contains unemployment rates for various countries in 1929 and 1932.

	1930	1931	1932	1933
U. rate	8.7	15.9	23.6	24.9
GDP growth	-8.5	-6.5	-12.9	-1.2

Table 1: Unemployment Rates and (real) GDP Growth during the Great Depression in the U.S. (source: thebalance.com)

The view that the current crisis is less severe because unemployment rates are smaller than during the Great Depression and so are the predicted declines in GDP is not uncommon. Somewhat implicit in this is that one compares current unemployment and GDP growth rates with the extremal numbers during the Great Depression. What is certainly a reason for considerable concern is the speed with which these changes occur today. For example, Australia's unemployment rate is expected to double in June from 5% to 10% while the unemployment rate in the U.S. has roughly quadrupled within four weeks from 4.4% in March to an expected 17% in April. It certainly already exceeds the unemployment rate in the U.S. in the first year of the Great Depression, that is, in 1930 (see Table 1). Figure 6 displays the dramatic increase in unemployment filings in the U.S. in the second half of March and the first half of April 2020. The precipitous decline in retail sales in the U.S. is displayed in Figure 7.

A major reason for the severity and length of the economic crisis that became the Great Depression was the slow and ultimately misguided policy reaction. Rather than increasing Government spending and money supply, the U.S. Government decreased both. While there are good reasons to believe that, today, these mistakes will be avoided, the possibility that, until there is a vaccine, the ongoing threat from the coronavirus renders these hard lessons learned useless. Swift policy responses lose their effectiveness if people cannot go back to work because of a looming public health disaster if they do.

Country	1929	1932
Australia	11.1	29.0
Austria	12.3	26.1
Belgium	4.3	39.7
Canada	5.7	22.0
Czechoslovakia	2.2	13.5
Denmark	15.5	31.7
Germany	n.a.	31.7
Japan	n.a.	6.8
Netherlands	7.1	29.5
Norway	15.4	30.8
Poland	4.9	11.8
Sweden	10.7	22.8
Switzerland	3.5	21.3
United Kingdom	10.4	22.1
United States	3.2	24.9
Average	8.2	25

Table 2: Unemployment Rates during the Great Depression (source: Reinhart and Rogoff (2009, p.269))

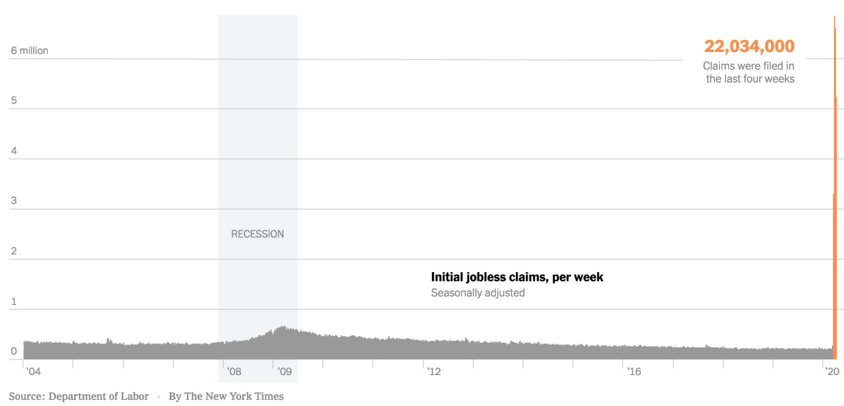

Figure 6: Weekly unemployment claims in the U.S. (Source: New York Times, 16 April 2020).

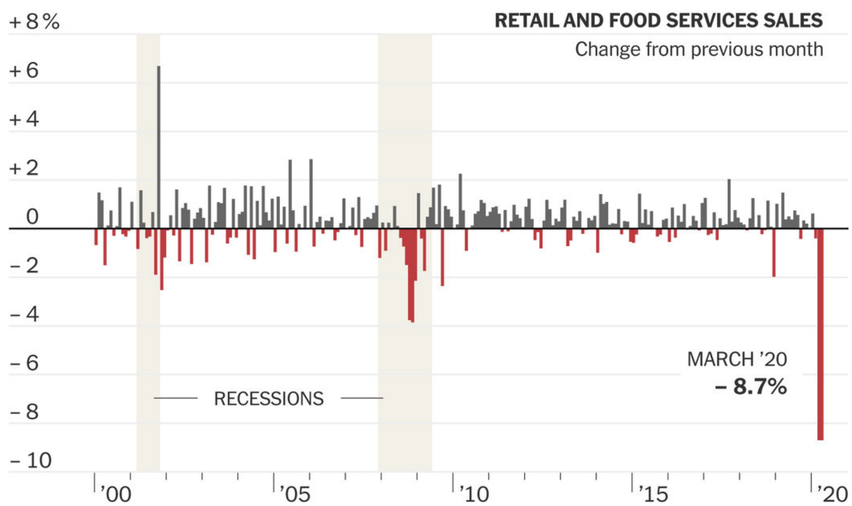

Figure 7: Retail sales. (Source: New York Times, 14 April 2020).