C A R F   W o r k i n g   P a p e r

CARF-F-620

Lockdown Policy Rules with a Hospital Capacity Constraint

Hongtao Li
The University of Tokyo

Taisuke Nakata
The University of Tokyo

Hiroki Sakamoto
The University of Tokyo

Hiroyuki Uneya
The University of Tokyo

March 3, 2026

CARF Working Papers can be downloaded without charge from:
https://www.carf.e.u-tokyo.ac.jp/research/

Lockdown Policy Rules with a Hospital Capacity Constraint[*]

Hongtao Li[†]   Taisuke Nakata[‡]   Hiroki Sakamoto[§]   Hiroyuki Uneya[¶]

March 3, 2026

Abstract

We analyze how hospital capacity affects health and economic outcomes during a pandemic using a macro-SIR model featuring a lockdown policy rule. The rule instructs the government to implement a lockdown when the number of ICU patients exceeds a trigger threshold—motivated broadly by a hospital capacity constraint—and to lift the lockdown when it falls below a lifting threshold. When vaccines are not available, we find that the government can reduce both COVID-19 deaths and economic loss by raising the trigger threshold in some situations. When the vaccine rollout is expected to begin in the near future, we find that the government can reduce both COVID-19 deaths and economic loss by lowering the trigger threshold.

JEL Codes: E17, I18, H12

Keywords: Epi-macro Model, Hospital Capacity, ICU, Lockdown, Macro-SIR Model, Pandemic, Vaccine.

[†]Graduate School of Public Policy, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan. Email address: li-hongtao@g.ecc.u-tokyo.ac.jp.

[‡]Graduate School of Economics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan. Email address: taisuke.nakata@e.u-tokyo.ac.jp.

[§]Department of Mathematical Informatics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan. Email address: soccer-books0329@g.ecc.u-tokyo.ac.jp.

[¶]Graduate School of Economics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan. Email address: arrow722679@g.ecc.u-tokyo.ac.jp.

1 Introduction

During the COVID-19 pandemic, governments often imposed a lockdown to contain the spread of infection. Lockdowns are often imposed for the purpose of preventing the number of severely-ill patients from exceeding the hospital capacity allocated to the disease. The government can adjust the hospital capacity allocated to the disease over the course of a pandemic by allocating hospital resources differently or by affecting the incentives of hospitals to accept COVID-19 patients. Such an adjustment in hospital capacity can affect health and economic outcomes during a pandemic by affecting the timing and duration of lockdowns.

We analyze how adjusting hospital capacity affects health and economic outcomes in a macro-SIR model with a lockdown policy rule. The policy rule instructs the government to implement a lockdown when the number of ICU patients exceeds a certain threshold—a trigger threshold—and to lift the lockdown when the number of ICU patients falls below a certain threshold—a lifting threshold. The trigger threshold is specified in terms of the number of ICU patients but is motivated more broadly by the constraints on hospital or medical capacity in general. Our main exercise is to analyze how the adjustment of the trigger threshold affects health and economic outcomes under various scenarios regarding vaccine rollout.

We consider the following four vaccine scenarios. In the first scenario, we assume that there is no vaccine available throughout the pandemic. In the second scenario, the vaccine rollout begins 10 weeks from the start of the pandemic. In the third scenario, the vaccine rollout begins 2 years from the start of the pandemic. In the fourth, final scenario, we assume that the vaccine rollout has ended, but that sufficiently many susceptible individuals are still left, and herd immunity has not been achieved. We call the first two scenarios the main scenarios and the last two scenarios the additional scenarios. We characterize the last two scenarios as "additional" because the results of the last two scenarios can be interpreted as a variation or a combination of the first two scenarios.

In the first main scenario without vaccines, we find that the government can reduce both COVID-19 deaths and economic loss by *raising* the trigger threshold in some situations. In this scenario, cumulative deaths decline non-monotonically as the trigger threshold is lowered, whereas average GDP loss increases almost monotonically. As a result, an increase in the trigger threshold can reduce both cumulative deaths and average GDP loss within a certain range. By contrast, in the second main scenario with an imminent vaccine rollout, we find that the government can reduce both COVID-19 deaths and economic loss by *lowering* the trigger threshold in some situations. In this scenario, the average GDP loss increases non-monotonically as the trigger threshold is lowered, whereas cumulative deaths decline

monotonically. As a result, lowering the trigger threshold can reduce both cumulative deaths and average GDP loss within a certain range.

In the first additional scenario in which the vaccine rollout starts two years from the beginning of the simulation, we find that the government can reduce both COVID-19 deaths and economic loss by raising the trigger threshold in some situations and by lowering the trigger threshold in other situations. This scenario can be interpreted as combining the mechanisms highlighted in the two main scenarios. As a result, two types of non-monotonicity responsible for creating the possibility of Pareto improvement both exist in this first additional scenario. On the other hand, the second additional scenario—one after vaccine rollout—is qualitatively the same as the first main scenario with no vaccine rollout. As a result, in the second additional scenario, the government can reduce both COVID-19 deaths and economic loss by raising the trigger threshold in some situations.

Our paper is related to the literature that analyzes optimal lockdown policies using macro-SIR models. Some authors analyze optimal lockdown policy in reduced-form macro-SIR models that abstract from explicit microfoundations. See, for example, Acemoglu et al. (2021), Alvarez et al. (2021), Aspri et al. (2021), Calvia et al. (2024), Hritonenko and Yatsenko (2024), and Prieur et al. (2024), among others. Other authors study optimal lockdown policy in micro-founded macro-SIR models. See, for example, Eichenbaum et al. (2021), Fajgelbaum et al. (2021), Farboodi et al. (2021), Glover et al. (2022), Goenka et al. (2024), Jones et al. (2021), among many others. Within this optimal-policy literature, some authors incorporate hospital or ICU capacity constraints in their analysis, either directly as a hard constraint (Acemoglu et al. (2021), Loertscher and Muir (2021), and Miclo et al. (2022)) or indirectly by specifying the fatality rate that increases with the number of infected or ICU individuals (Casares et al. (2025), Caulkins et al. (2021), Garriga et al. (2022), and Ma et al. (2022)). We differ from all of these papers because we analyze lockdown policy rules—instead of the optimal policy. When the policy is chosen optimally, imposing a hospital capacity constraint is typically welfare-reducing. However, with lockdown policy rules, we find situations where imposing a hospital capacity constraint can improve both health and economic outcomes, thus enhancing welfare.

Our paper is closely related to a set of papers analyzing how lockdown policy *rules* affect health and economic outcomes in macro-SIR models.[1] Bar-On et al. (2023) study a time-based cyclic rule that alternates between opening and lockdown according to a predetermined calendar. Fukao and Shioji (2022) analyze a health-economy tradeoff in a macro-SIR model

---

[1]Some papers analyze lockdown policy rules using pure epidemiological models without economic considerations. See, for example, Davies et al. (2020), Della Rossa et al. (2020), Kissler et al. (2020), and Yang et al. (2021), among many others.

featuring a policy rule in which lockdown intensity increases with the number of infected individuals. Kubota (2021) analyzes health-economy tradeoffs in a macro-SIR model featuring an ICU targeting policy rule that keeps the number of ICU patients roughly constant. We differ from these three papers because we focus on ICU-triggered lockdown rules with trigger and lifting thresholds.

Our paper complements Fujii and Nakata (2021). Fujii and Nakata (2021) assume a trigger-based lockdown policy rule in a macro-SIR model similar to ours. They analyze how the lifting—instead of trigger—threshold affects health and economic outcomes and find that lowering the lifting threshold can improve both health and economic outcomes when vaccine rollout is expected to begin in the near future. We differ from Fujii and Nakata (2021) because we analyze the implications of adjusting the trigger threshold and consider various assumptions regarding the timing of vaccine rollout.

Some authors have analyzed the effects of lockdown policies on health and economic outcomes by specifying an exogenous path of lockdown intensities, as opposed to assuming the optimal policy or a policy rule. Examples include Ascari et al. (2023), Baqaee et al. (2020), Fu et al. (2020), Hosono (2021), and Kubota (2021), among others. The assumption of an exogenous path is particularly useful for policy analyses using empirically realistic macro-SIR models. We complement these studies by analyzing health-economy tradeoffs using a lockdown policy rule.

This paper is organized as follows. Section 2 presents the model, the vaccine rollout assumptions, and the parameterization. Section 3 presents the simulation results in the absence of vaccine rollout. Section 4 presents the simulation results under the impending vaccine rollout. Section 5 presents the simulation results under additional vaccine rollouts. Section 6 concludes.

2   Model

2.1   SIR Model

Following Acemoglu et al. (2021) and Alvarez et al. (2021), we use a reduced-form macro-SIR model for our analysis. The model is formulated in discrete time, where each period corresponds to one week. The following equations describe the dynamics of the model's epidemic part:

$$S_{t+1} - S_t = -N_t^S - \frac{S_t}{(S_t + R_t)} V_t \text{Pop}, \tag{1}$$

$$I_{t+1} - I_t = N_t^S - \gamma I_t - \delta I_t, \tag{2}$$

$$R_{t+1} - R_t = \gamma I_t - \frac{R_t}{(S_t + R_t)} V_t \text{Pop}, \tag{3}$$

$$X_{t+1} - X_t = -N_t^X + \frac{S_t}{(S_t + R_t)} V_t \text{Pop}, \tag{4}$$

$$Y_{t+1} - Y_t = N_t^X - \gamma Y_t - (1 - e)\delta Y_t, \tag{5}$$

$$Z_{t+1} - Z_t = \gamma Y_t + \frac{R_t}{(S_t + R_t)} V_t \text{Pop}, \tag{6}$$

$$D_{t+1} - D_t = \delta I_t + (1 - e)\delta Y_t, \tag{7}$$

$$S_t + I_t + R_t + X_t + Y_t + Z_t + D_t = \text{Pop}. \tag{8}$$

S_t, I_t, R_t, and D_t denote the numbers of unvaccinated susceptible, infected, recovered individuals, and cumulative deaths at week t, respectively. X_t, Y_t, and Z_t denote the number of vaccinated susceptible, infected, and recovered individuals, respectively. V_t denotes the proportion of newly vaccinated individuals to the population at week t. The parameter e denotes the effectiveness of vaccination in reducing infection, ICU, and mortality rates. N_t^S and N_t^X denote new infections among unvaccinated and vaccinated susceptible individuals, respectively. Parameters γ, ϕ, and δ denote recovery, ICU, and mortality rates, respectively.

We assume that vaccination targets individuals who are unvaccinated and not currently infected and does not distinguish between susceptible and recovered individuals. In each week t, the government vaccinates V_tPop individuals drawn from $S_t + R_t$ in proportion to their population shares: a fraction $S_t/(S_t + R_t)$ comes from S_t and the remaining fraction $R_t/(S_t + R_t)$ comes from R_t. Newly vaccinated individuals drawn from S_t move to X_{t+1}, while those drawn from R_t move to Z_{t+1}.

$$N_t^S = \beta (1 - h\alpha_t)^2 (I_t + Y_t) \frac{S_t}{\text{Pop}}, \tag{9}$$

$$N_t^X = (1 - e)\beta (1 - h\alpha_t)^2 (I_t + Y_t) \frac{X_t}{\text{Pop}}, \tag{10}$$

Newly infected cases from S_t (N_t^S) are determined by the matching between S_t and aggregate current infections ($I_t + Y_t$) (as shown in equation (9)). Newly infected cases from X_t (N_t^X) are determined by the matching between X_t and aggregate current infections ($I_t + Y_t$), adjusted for vaccine efficacy e (as shown in equation (10)). The term $\beta (1 - h\alpha_t)^2$ describes the mitigated transmission rate, which captures the relationship between the COVID-19 transmission and economic activities. The parameter β is the unmitigated transmission rate. The variable α_t denotes lockdown intensity—the degree of reduction in economic activities.

The parameter h represents the elasticity of the transmission rate to lockdown intensity. When a lockdown is in place, fewer new infections occur.

$$ICU_{t+1}^S - ICU_t^S = \phi N_t^S - \gamma ICU_t^S - \delta ICU_t^S, \tag{11}$$

$$ICU_{t+1}^X - ICU_t^X = (1-e)\phi N_t^X - \gamma ICU_t^X - (1-e)\delta ICU_t^X, \tag{12}$$

$$ICU_t = ICU_t^S + ICU_t^X. \tag{13}$$

ICU_t^S and ICU_t^X represent the numbers of unvaccinated and vaccinated ICU patients, respectively. Without loss of generality, we assume that the new ICU patients are a proportion of new infections. New ICU admissions among vaccinated individuals are scaled down by vaccine efficacy. ICU_t^S is a subset of I_t while ICU_t^X is a subset of Y_t. ICU_t denotes the aggregate ICU patients. Pop represents the entire population.

2.2 Lockdown Policy Rule and Output

We assume that the government chooses lockdown intensity at each period according to the following policy rule. When the number of ICU patients exceeds $\overline{ICU}$ (the trigger threshold), the government imposes a lockdown of intensity $c > 0$. When the number of ICU patients falls below $\underline{ICU}$ (the lifting threshold), the government lifts the lockdown. We interpret $\overline{ICU}$ as the effective ICU capacity allocated to COVID-19.

The lockdown intensity α_t under this lockdown policy rule is given by

$$\alpha_t = \begin{cases} c, & ICU_t \geq \overline{ICU}, \\ 0, & ICU_t < \underline{ICU}, \\ \alpha_{t-1}, & \text{Otherwise.} \end{cases} \tag{14}$$

and the lockdown intensity at time one is zero ($\alpha_1 = 0$).

The focus of our paper is how the choice of the trigger threshold—$\overline{ICU}$—affects epidemic and economic dynamics, paying particular attention to how the effects depend on the vaccine rollout. See Fujii and Nakata (2021) for the analysis of $\underline{ICU}$ under a specific scenario regarding vaccine rollout.

We assume that the output (GDP) per capita of the economy y_t is given by:

$$y_t = (1-\alpha_t)\frac{(S_t + X_t + \chi(I_t + Y_t) + R_t + Z_t)}{\text{Pop}}, \tag{15}$$

where χ is the productivity of the infected population. The output loss is determined by

lockdown intensity, lockdown duration, decline in productivity of the infected population, and cumulative deaths.

Average GDP loss is given by:

$$\frac{1}{T}\sum_{t=1}^{T}(1-y_t),\tag{16}$$

where T is the simulation horizon. We normalize GDP per capita before the pandemic to 1.

2.3 Parameterization

We parameterize the macro-SIR model at a weekly frequency. We set the recovery rate (γ) to 7/12, implying an average infectious period of 12 days, consistent with Kamo et al. (2022). We set the ICU hospitalization rate (ϕ) to 0.004 as in Zardini et al. (2021). We also set the mortality rate (δ) to 0.002 so that the implied infection fatality risk $\delta/(\gamma + \delta)$ (based on $\gamma = 7/12$) matches the population-averaged infection fatality risk reported by Zhang and Nishiura (2023). We set vaccine efficacy (e) to 0.7, consistent with estimates that a first dose of mRNA vaccines reduces infection by about 70% and severe outcomes (hospital demand or death) by about 70% in a large population-based study Chung et al. (2021). We set the unmitigated transmission rate (β) to 0.89, following the basic reproductive number (1.5) in Pindyck (2020). We set lockdown intensity (c) to 0.3, approximating the annualized 27.8% reduction in Japan's real GDP in 2020Q2 in Morita and Ono (2024). We set the elasticity of the transmission rate to lockdown intensity (h) to 1.5, which is consistent with that in Fujii and Nakata (2021). We set the productivity of the infected population (χ) to 0.8 as in Eichenbaum et al. (2021). We set the population (Pop) to 13,960,000, roughly matching the population of Tokyo.

Table 1: Parameters

Parameter	Description	Value	Reference
γ	Recovery rate	$\frac{7}{12}$	Kamo et al. (2022)
ϕ	ICU hospitalization rate	0.004	Zardini et al. (2021)
δ	Mortality rate	0.002	Zhang and Nishiura (2023)
e	Vaccine efficacy	0.7	Chung et al. (2021)
β	Unmitigated transmission rate	[0.89, 1.78]	Pindyck (2020)
c	Lockdown intensity	0.3	Morita and Ono (2024)
h	Elasticity of the transmission rate to lockdown intensity	1.5	Fujii and Nakata (2021)
χ	Productivity of the infected population	0.8	Eichenbaum et al. (2021)
Pop	Population in Tokyo	13,960,000	Tokyo Government[1]

[1] https://www.koho.metro.tokyo.lg.jp/2021/02/index.html.

Table 2: Initial Condition for Four Scenarios

Variable	Main Scenario 1	Main Scenario 2	Additional Scenario 1	Additional Scenario 2
S_1	$\text{Pop} - I_1 - R_1 - D_1$	$\text{Pop} - I_1 - R_1 - D_1$	$\text{Pop} - I_1 - R_1 - D_1$	$0.3 * (\text{Pop} - I_1 - R_1 - D_1)$
I_1	10,000	10,000	10,000	3,000
R_1	0	0	0	0
D_1	0	0	0	0
ICU_1^S	0	0	0	0
X_1	0	0	0	$0.7 * (\text{Pop} - I_1 - R_1 - D_1)$
Y_1	0	0	0	7,000
Z_1	0	0	0	0
ICU_1^X	0	0	0	0

Note: There is no vaccine rollout in the first main scenario; Vaccine rollout begins 10 weeks after the initial week in the second main scenario; Vaccine rollout begins 2 years after the initial week in the first additional scenario; Vaccine rollout begins before the initial week in the second additional scenario. All quantities are measured in persons.

2.4 Vaccine Rollout Scenarios

We simulate the model across four scenarios: two main scenarios and two additional scenarios. The simulation horizon (T) is 520 weeks (10 years). The lifting threshold ($\underline{ICU}$) is set to 30 across all scenarios. Table 2 shows that initial conditions are the same across all scenarios except for the second additional scenario.

In the first main scenario, we assume no vaccine rollout. We vary the trigger threshold for lockdown from 100 to 4,000 in increments of 1. The initial unvaccinated group and deaths $\{S_1, I_1, R_1, ICU_1^S, D_1\}$ are set to $\{\text{Pop} - I_1 - R_1 - D_1, 10000, 0, 0, 0\}$. The initial vaccinated group $\{X_1, Y_1, Z_1, ICU_1^X\}$ is set to $\{0, 0, 0, 0\}$.

In the second main scenario, we assume that the vaccine rollout begins 10 weeks after the initial week.[2] To simplify the analysis, and without loss of generality, we assume that vaccine rollout is completed in one week. We assume that 70% of the population will be vaccinated.[3] We adjust the trigger threshold of lockdown from 40 to 1,000 in increments of 1. The initial conditions in this scenario are the same as those in the first main scenario.

In the first additional scenario, we assume that vaccine rollout begins 104 weeks (two years) after the initial week. Without loss of generality, we again assume that the vaccine rollout will be completed in one week. We assume that 70% of the population will be vaccinated, as in the second main scenario. We adjust the trigger threshold of lockdown from 100 to 4,000 in increments of 1. The initial conditions in this scenario are the same as those in the first main scenario.

In the second additional scenario, we assume that vaccine rollout begins before the ini-

---

[2]The exact rollout date is not essential; what matters is that vaccination begins in the near future. The main takeaways from the analysis are robust to alternative rollout dates, such as 20, 30, 40 weeks, etc.

[3]Allowing vaccination to proceed gradually at a constant rate over multiple weeks does not alter the key takeaways from the analysis.

tial week, and 70% of the population has been vaccinated. In this scenario, we set β to 1.78, corresponding to a basic reproductive number of 3. If we set β to 0.89, as in the other scenarios, there would be no additional infection wave after vaccine rollout. However, in reality, many countries experienced further infection waves even after vaccines were distributed to most citizens. Increasing the base transmission rate in this scenario allows us to capture such situations. We adjust the trigger threshold of lockdown from 100 to 2,500 in increments of 1. Initial unvaccinated susceptible and infected individuals $\{S_1, I_1\}$ are set to $\{0.3 * (\text{Pop} - I_1 - R_1 - D_1), 3000\}$ and initial vaccinated susceptible and infected individuals $\{X_1, Y_1\}$ are set to $\{0.7 * (\text{Pop} - I_1 - R_1 - D_1), 7000\}$, while initial conditions of other groups are the same as in previous scenarios.
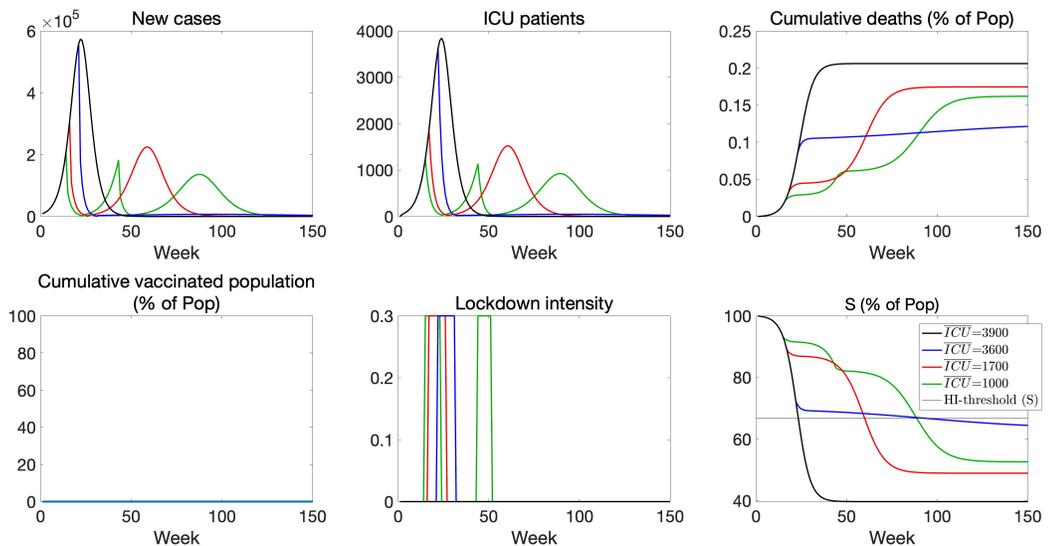
3 Results: Main Scenario 1

Figure 1: Model Dynamics in the No-Vaccine Scenario

Notes: The solid black, blue, red, and green lines represent the dynamics under trigger thresholds of 3,900, 3,600, 1,700, and 1,000 ICU patients, respectively.

Figure 1 shows the epidemic dynamics in the first main scenario—a scenario without a vaccine. Four lines in each panel correspond to four trigger thresholds. When the trigger threshold is set at 3,900 ICU patients (black lines), the number of ICU patients never exceeds the threshold, and no lockdowns are imposed. In this case, there is a single infection wave. When the threshold is lower, the number of ICU patients exceeds the threshold at some point and triggers a lockdown. When the threshold is 3,600 (blue lines) and 1,700 (red lines), there

8

is a single lockdown with two infection waves. When the threshold is further lowered to 1,000 (green lines), we observe two lockdowns and three infection waves. In general, a lower trigger threshold can lead to more frequent lockdowns and thus greater economic loss.
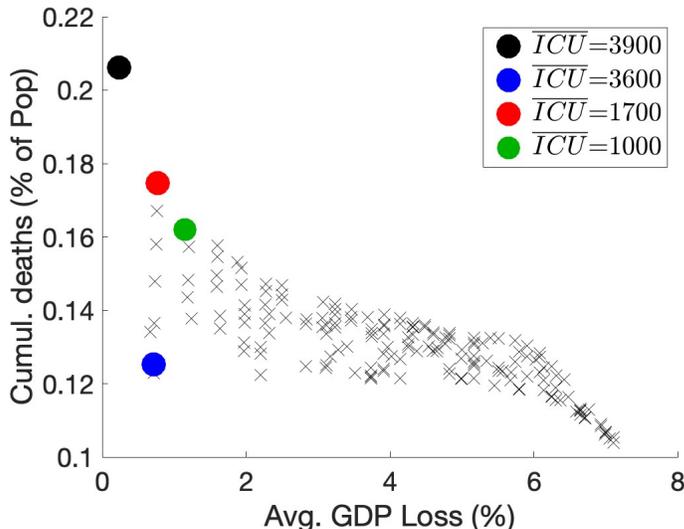

Figure 2: Cumulative Deaths and Average GDP Loss in the No-Vaccine Scenario

Notes: This figure shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds in the no-vaccine scenario. The black, blue, red, and green dots represent the trigger thresholds of 3,900, 3,600, 1,700, and 1,000 ICU patients, respectively. The simulation horizon is 520 weeks.

Figure 2 shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds in this scenario. Each cross corresponds to the outcomes associated with a particular trigger threshold. Four colored dots in Figure 2 correspond to four lines of the same color in Figure 1. The cloud of crosses indicates that there is a global tradeoff—that is, a tradeoff over the full range of thresholds—between cumulative deaths and average GDP loss—thresholds associated with fewer deaths tend to be associated with larger average GDP losses. However, locally—that is, within a narrower range of thresholds—we observe that some thresholds (e.g., 3,600, the blue dot) lead to both lower cumulative deaths and smaller average GDP loss than others (e.g., 1,000, the green dot). That is, within some local ranges, the government can reduce both cumulative deaths and average GDP loss at the same time by raising the trigger threshold.

The possibility of such Pareto improvement arises because cumulative deaths decline non-monotonically as the threshold is lowered, whereas average GDP loss increases almost monotonically. These patterns can be seen clearly in Figure 3, which shows how the trigger threshold affects cumulative deaths and average GDP loss separately. According to Figure 3a, as the government lowers the trigger threshold, cumulative deaths decline, but non-

9

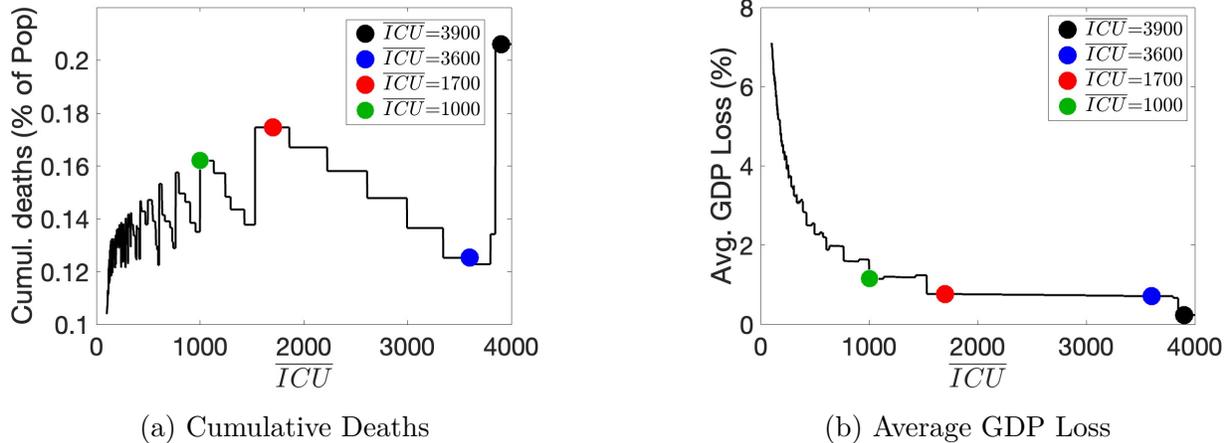(a) Cumulative Deaths   (b) Average GDP Loss

Figure 3: The Effects of Adjusting the Trigger Threshold in the No-Vaccine Scenario

Notes: This figure shows the relationship between trigger thresholds and two outcomes in the no-vaccine scenario. In Panel (a), the black line shows how cumulative deaths vary with the trigger threshold. In Panel (b), the black line shows how the average GDP loss varies with the trigger threshold. The black, blue, red, and green dots represent the trigger thresholds of 3,900, 3,600, 1,700, and 1,000 ICU patients, respectively. The simulation horizon is 520 weeks.

monotonically. On the other hand, according to Figure 3b, as the government lowers the trigger threshold, the average GDP loss increases almost monotonically. These patterns imply that, within specific ranges, a higher trigger threshold can result in both lower cumulative deaths and smaller average GDP loss than a lower one.

Phase diagrams in Figures 4a and 4b help us understand why the cumulative number of deaths changes with the trigger threshold in a non-monotonic way. Figure 4a shows the phase diagram for four different trigger thresholds associated with different numbers of lockdowns. When the trigger threshold is sufficiently high (e.g., 3,900, the black dot), there is no lockdown. As the trigger threshold is lowered, lockdowns become more frequent. More frequent lockdowns are associated with a larger terminal value of S—that is S closer to the herd immunity threshold—which in turn implies lower cumulative deaths, as pointed out by Moll (2020). That is, a lower trigger threshold is associated with fewer deaths when it induces a larger number of lockdowns.

On the other hand, conditional on the number of lockdowns remaining constant, a lower trigger threshold is associated with a higher cumulative death toll. This relationship can be seen in Figure 4b, which shows the phase diagram for four different trigger thresholds associated with a single lockdown. According to the figure, as the trigger threshold is lowered, the final S declines. The smaller final S means that more people have been infected—and thus more people have died—before the pandemic ends.

To summarize, lowering the trigger threshold reduces the cumulative deaths globally. However, conditional on the number of lockdowns remaining unchanged, lowering the trig-
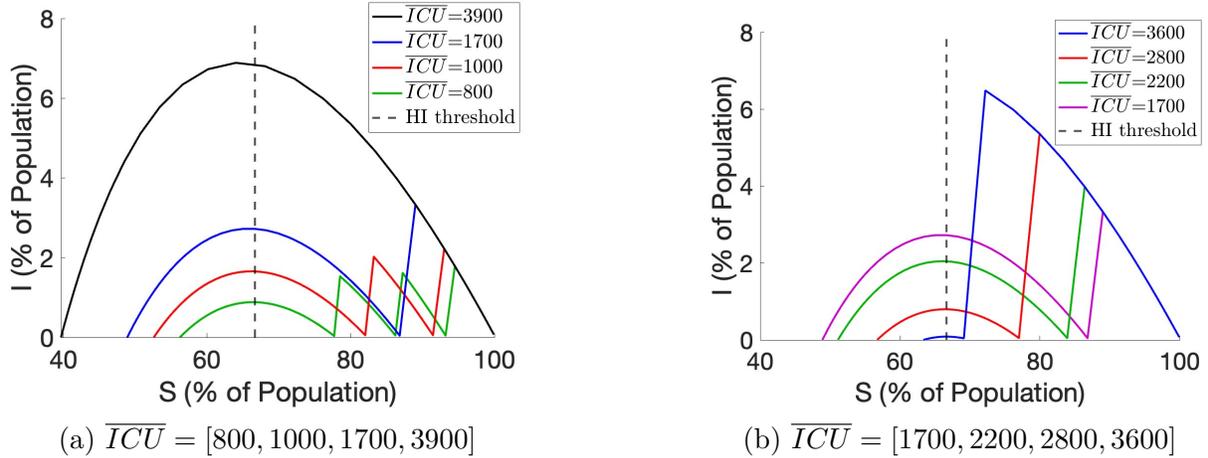
(a) $\overline{ICU} = [800, 1000, 1700, 3900]$ (b) $\overline{ICU} = [1700, 2200, 2800, 3600]$

Figure 4: Phase Diagrams in the No-Vaccine Scenario

Notes: These two phase diagrams show the relationship between S and I, indicating how the final S is associated with a trigger threshold. Panel (a) shows the phase diagram for four trigger thresholds associated with different numbers of lockdowns. The black, blue, red, and green lines represent the trigger thresholds of 3,900, 1,700, 1,000, and 800 ICU patients, respectively. Panel (b) shows the phase diagram for four trigger thresholds associated with a single lockdown. The blue, red, green, and purple lines represent the trigger thresholds of 3,600, 2,800, 2,200, and 1,700 ICU patients, respectively. The simulation horizon is 520 weeks.

ger threshold can increase the cumulative deaths locally, leading to the non-monotonicity observed in Figure 3a. This non-monotonicity in turn creates the possibility of Pareto improvement in health and economic outcomes.

4   Results: Main Scenario 2

Figure 5 shows the epidemic dynamics in the second main scenario, in which vaccine rollout begins 10 weeks after the initial week. Four lines in each panel show the dynamics under four alternative trigger thresholds. When the trigger threshold is set at 800 ICU patients (black lines), the number of ICU patients never exceeds the threshold. In this case, no lockdowns are imposed, and there is a single infection wave. As the threshold is lowered to 200 (blue lines) and 100 (red lines), the number of ICU patients exceeds the threshold, triggering a lockdown. We observe two infection waves in this case. When the threshold is further lowered to 50 (green lines), it triggers three lockdowns, and there are four infection waves. As in the previous section, a lower trigger threshold can also lead to more frequent lockdowns.
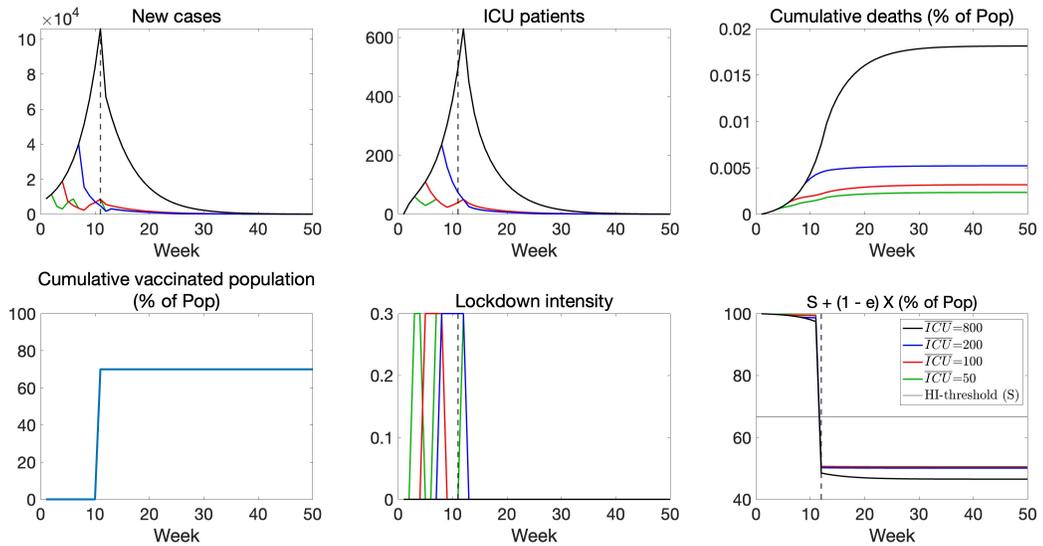
11

Figure 5: Model Dynamics in the Impending-Vaccine Scenario

Notes: The solid black, blue, red, and green lines represent the dynamics under trigger thresholds of 800, 200, 100, and 50 ICU patients, respectively. The vaccine rollout begins 10 weeks after the initial week, as shown in the bottom-left panel. The vertical, dashed black line represents the time of vaccine rollout.

Figure 6 shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds in this scenario. Each cross represents the outcomes associated with a particular trigger threshold. Four colored dots in Figure 6 correspond to four lines of the same color in Figure 5. Across the set of crosses, we again observe a global tradeoff between cumulative deaths and average GDP loss—thresholds associated with fewer deaths tend to be associated with larger average GDP losses. However, locally, we observe that some thresholds (e.g., 100, the red diamond marker) lead to lower cumulative deaths and smaller average GDP loss than others (e.g., 200, the blue diamond marker). That is, within some local ranges, the government can reduce both cumulative deaths and average GDP loss at the same time by lowering the trigger threshold.
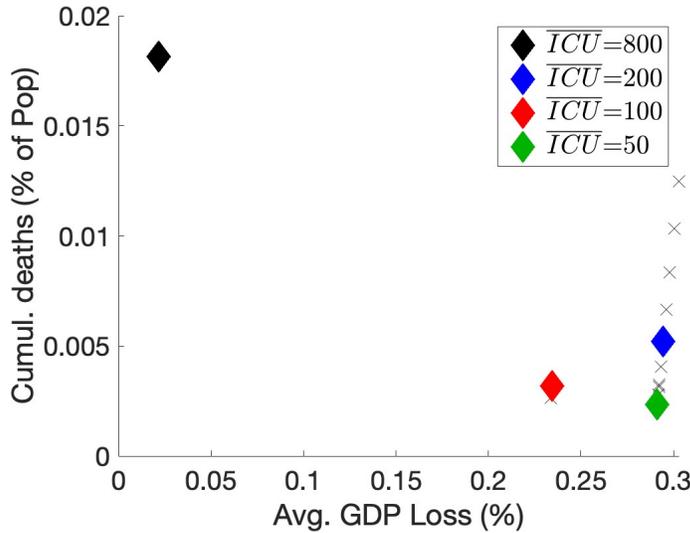
Figure 6: Cumulative Deaths and Average GDP Loss in the Impending-Vaccine Scenario

Notes: This figure shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds in the impending-vaccine scenario. The black, blue, red, and green diamond markers represent the trigger thresholds of 800, 200, 100, and 50 ICU patients, respectively. The simulation horizon is 520 weeks.

The possibility of such Pareto improvement arises because the average GDP loss increases non-monotonically as the threshold is lowered, whereas cumulative deaths decline monotonically. These patterns can be seen in Figure 7, which shows how the trigger threshold affects cumulative deaths and average GDP loss separately. According to Figure 7a, as the government lowers the trigger threshold, cumulative deaths decline monotonically. On the other hand, according to Figure 7b, as the government lowers the trigger threshold, the average GDP loss increases, but non-monotonically. These patterns imply that, over certain ranges, a lower trigger threshold (e.g., 100, the red diamond marker) can result in lower cumulative deaths and smaller average GDP loss simultaneously than a higher one (e.g., 200, the blue diamond marker).

Figures 8a and 8b help us understand why the average GDP loss changes with the trigger threshold in a non-monotonic way. Both panels show the dynamics of cumulative GDP loss under a small number of selected trigger thresholds. In Figure 8a, the selected trigger thresholds correspond to different numbers of lockdowns. In Figure 8b, selected trigger thresholds are associated with a single lockdown.

According to Figure 8a, when the trigger threshold is sufficiently high (e.g., 800, the black line), no lockdowns are imposed. As the government lowers the trigger threshold, lockdowns become more frequent. More frequent lockdowns are associated with a longer aggregate lockdown duration, which in turn implies greater cumulative GDP loss. On the other hand, conditional on the number of lockdowns remaining constant, a lower trigger threshold is
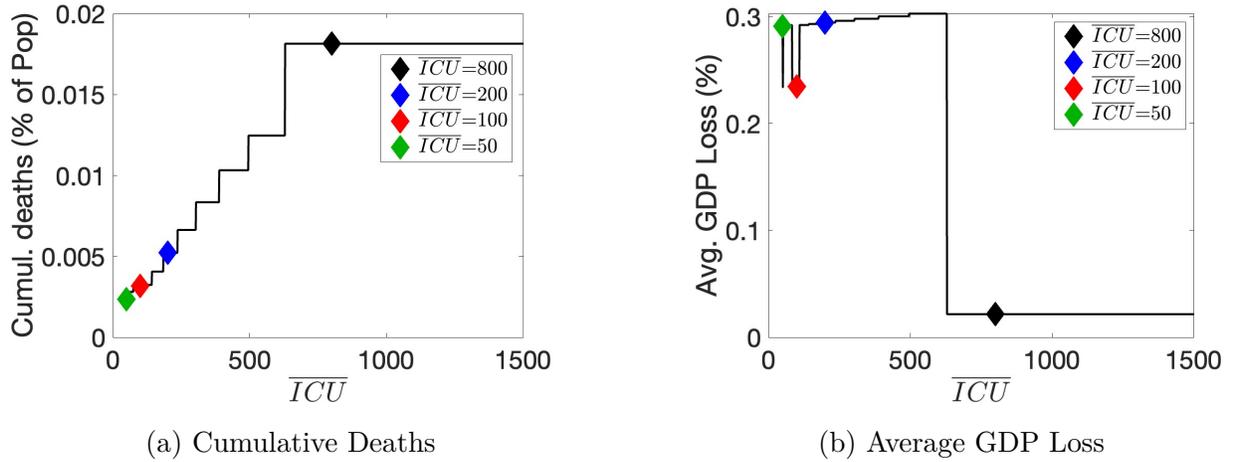
13

(a) Cumulative Deaths

(b) Average GDP Loss

Figure 7: The Effects of Adjusting the Trigger Threshold in the Impending-Vaccine Scenario

Notes: This figure shows the relationship between trigger thresholds and two outcomes in the impending-vaccine scenario. In Panel (a), the black line shows how cumulative deaths vary with the trigger threshold. In Panel (b), the black line shows how the average GDP loss varies with the trigger threshold. The black, blue, red, and green diamond markers represent the trigger thresholds of 800, 200, 100, and 50 ICU patients, respectively. The simulation horizon is 520 weeks.

(a) $\overline{ICU} = [50, 100, 800]$
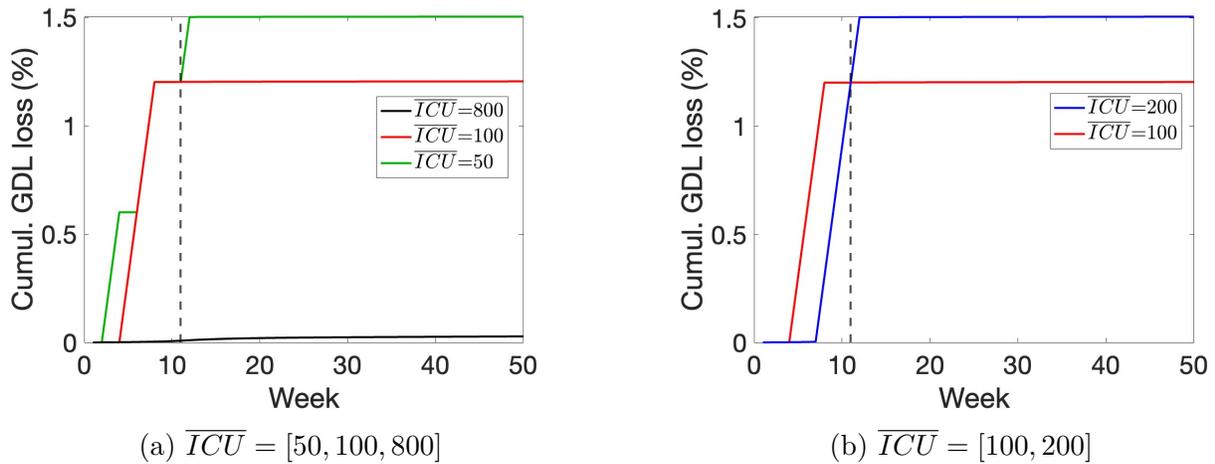
(b) $\overline{ICU} = [100, 200]$

Figure 8: Dynamics of Cumulative GDP Loss with Alternative Trigger Thresholds

Notes: These two panels show how the dynamics of cumulative GDP loss vary with the trigger threshold. Panel (a) shows the dynamics for three trigger thresholds associated with different numbers of lockdowns. Panel (b) shows the dynamics for two trigger thresholds associated with a single lockdown. The black, blue, red, and green lines represent the trigger thresholds of 800, 200, 100, and 50 ICU patients, respectively. The simulation horizon is 520 weeks.

14

associated with smaller average GDP loss. According to Figure 8b, as the government lowers the trigger threshold, the lockdown starts earlier. Because the lockdown starts earlier, the numbers of infections and ICU patients are lower at the onset of the lockdown. It takes less time for the number of ICU patients to decline below the lifting threshold, and thus the lockdown ends earlier, implying a shorter lockdown duration. A shorter lockdown duration implies a smaller cumulative GDP loss and thus a smaller average GDP loss.[4]

To summarize, lowering the trigger threshold raises the average GDP loss globally. However, lowering the trigger threshold can reduce the average GDP loss locally—conditional on the number of lockdowns remaining unchanged—leading to the non-monotonicity observed in Figure 7b. This non-monotonicity creates the possibility of Pareto improvement in health and economic outcomes.

5   Results: Additional Scenarios

To better understand how the ICU capacity constraint affects health and economic outcomes and how that relationship depends on vaccine assumptions, we consider two additional scenarios: (i) vaccine rollout begins in the intermediate future, and (ii) vaccine rollout has ended by the beginning of the simulation.

5.1   Additional Scenario 1

In this scenario, we assume that the vaccine rollout begins 104 weeks (2 years) after the beginning of the pandemic. Figure 9 shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds in this scenario. Each cross represents the outcomes associated with a particular trigger threshold. Globally, there is a tradeoff between cumulative deaths and average GDP loss, as in the two main scenarios. Locally, we observe that some thresholds result in fewer cumulative deaths and a smaller average GDP loss at the same time.

We observe that the government can reduce both cumulative deaths and average GDP loss by raising the trigger threshold in some situations and by lowering the trigger threshold in other situations. In Figure 9, raising the trigger threshold from 1,000 (the green dot) to 3,400 (the blue dot) reduces both cumulative deaths and average GDP loss; lowering the trigger threshold from 700 (the blue diamond marker) to 480 (the red diamond marker) reduces both cumulative deaths and average GDP loss.

---

[4]The cumulative GDP loss will not increase after week 50. Therefore, a smaller cumulative GDP loss implies a smaller average GDP loss.
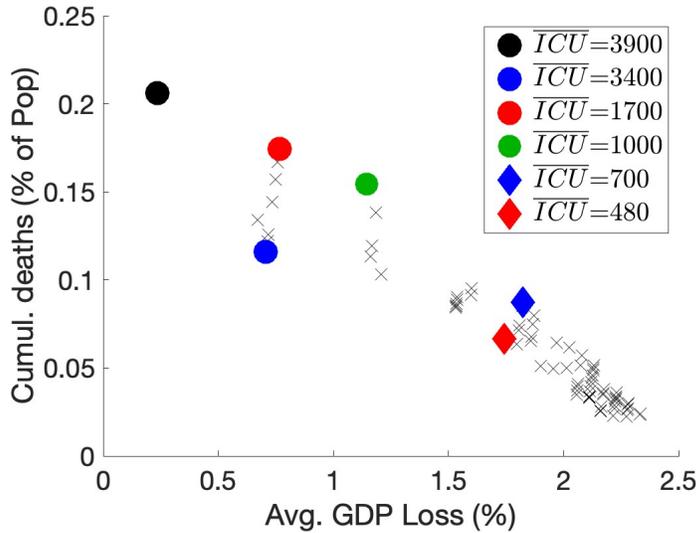
Figure 9: Cumulative Deaths and Average GDP Loss under an Intermediate-term Vaccine Scenario

Notes: The figure shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds under the intermediate-term vaccine rollout. The black, blue, red, and green dots represent the outcomes associated with trigger thresholds of 3,900, 3,400, 1,700, and 1,000 ICU patients. The blue and red diamond markers represent 700 and 480 ICU patients, respectively. The simulation horizon is 520 weeks.

This result is intuitive because an intermediate-term vaccine rollout combines the two mechanisms highlighted in the main scenarios. For relatively high trigger thresholds, the outcomes are qualitatively similar to those in the no-vaccine scenario because herd immunity is reached before the vaccine rollout; therefore, the vaccine rollout occurs after the epidemic is largely over and has a limited effect on subsequent infections and thus deaths. For relatively low trigger thresholds, outcomes are qualitatively similar to those in the impending-vaccine rollout scenario because herd immunity is reached only after the vaccine rollout; therefore, the vaccine rollout occurs while the epidemic is still ongoing and has a large effect on subsequent infections and thus deaths.

5.2 Additional Scenario 2

In this scenario, we consider a post-vaccine-rollout environment. Figure 10 shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds in this scenario. Each cross represents the outcomes associated with a particular trigger threshold. Across the set of crosses, we observe a global tradeoff between cumulative deaths and average GDP loss—thresholds associated with lower cumulative deaths tend to be associated with higher average GDP loss. However, locally, we observe that some thresholds result in fewer cumulative deaths and smaller average GDP loss simultaneously. For example, a lower

trigger threshold of 600 (the green dot) is associated with more cumulative deaths and a higher average GDP loss than a higher trigger threshold of 1,900 (the blue dot).

This result is qualitatively similar to that in the first main scenario (no-vaccine scenario). In both this scenario and the first main scenario, no further vaccine rollout occurs. The only difference between the two scenarios is the initial level of susceptible population. Both scenarios feature the two key forces responsible for creating the possibility of Pareto improvement—a non-monotonic relationship between the trigger threshold and cumulative deaths and a monotonic relationship between the trigger threshold and average GDP loss—highlighted earlier in Figure 3a and Figure 3b.
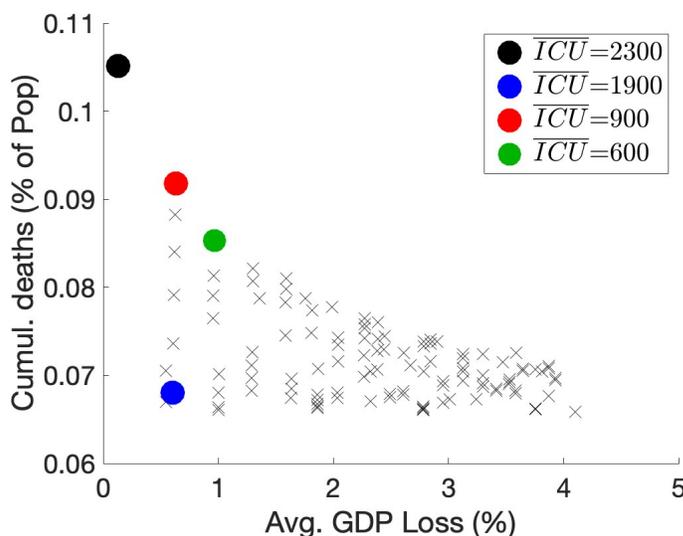

Figure 10: Cumulative Deaths and Average GDP Loss in the Post-Vaccine Scenario

Notes: This figure shows pairs of cumulative deaths and average GDP loss associated with various trigger thresholds in the post-vaccine-rollout scenario. The black, blue, red, and green dots represent the trigger thresholds of 2,300, 1,900, 900, and 600 ICU patients, respectively. The simulation horizon is 520 weeks.

6 Conclusion

We have analyzed how the choice of hospital capacity constraint affects health and economic outcomes during a pandemic using a macro-SIR model featuring a lockdown policy rule. We find that the government can reduce both COVID-19 deaths and economic loss by (i) raising the lockdown trigger threshold in the absence of vaccine rollout, and (ii) lowering the lockdown trigger threshold when vaccine rollout is imminent. The key mechanism generating the first result is a non-monotonic relationship between the trigger threshold and cumulative deaths, whereas the key mechanism generating the second result is a non-monotonic relationship between the trigger threshold and GDP loss.

During a pandemic, governments must decide how much medical or hospital capacity to allocate to the infectious disease. Because hospital capacity affects lockdown decisions, it affects not only infection outcomes but also economic outcomes. Our research highlights the complex relationship between hospital capacity and the health-economy tradeoff. Further research would help sharpen our understanding of the role of hospital capacity in managing health-economic tradeoffs during a pandemic.

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2021). Optimal targeted lockdowns in a multigroup SIR model. *American Economic Review: Insights 3*(4), 487–502.

Alvarez, F., D. Argente, and F. Lippi (2021). A simple planning problem for COVID-19 lock-down, testing, and tracing. *American Economic Review: Insights 3*(3), 367–82.

Ascari, G., A. Colciago, and R. Silvestrini (2023, jul). Business dynamism, sectoral reallocation and productivity in a pandemic. *European Economic Review 156*, 104473.

Aspri, A., E. Beretta, A. Gandolfi, and E. Wasmer (2021). Mortality containment vs. economics opening: Optimal policies in a SEIARD model. *Journal of Mathematical Economics 93*, 102490.

Baqaee, D. R., E. Farhi, M. J. Mina, and J. H. Stock (2020). Reopening scenarios. NBER Working Paper 27244, National Bureau of Economic Research.

Bar-On, Y., T. Baron, O. Cornfeld, and E. Yashiv (2023). When to lock, not whom: Managing epidemics using time-based restrictions. *Review of Economic Dynamics 51*, 292–321.

Calvia, A., F. Gozzi, F. Lippi, and G. Zanco (2024). A simple planning problem for COVID-19 lockdown: a dynamic programming approach. *Economic Theory 77*(1), 169–196.

Casares, M., P. Gomme, and H. Khan (2025). Private versus social responses to a pandemic. *Journal of Economic Dynamics and Control 181*, 105202.

Caulkins, J. P., D. Grass, G. Feichtinger, R. F. Hartl, P. M. Kort, A. Prskawetz, A. Seidl, and S. Wrzaczek (2021). The optimal lockdown intensity for COVID-19. *Journal of Mathematical Economics 93*, 102489.

Chung, H., S. He, S. Nasreen, M. E. Sundaram, S. A. Buchan, S. E. Wilson, B. Chen, A. Calzavara, D. B. Fell, P. C. Austin, K. Wilson, K. L. Schwartz, K. A. Brown, J. B. Gubbay, N. E. Basta, S. M. Mahmud, C. H. Righolt, L. W. Svenson, S. E. MacDonald, N. Z. Janjua, M. Tadrous, J. C. Kwong, and Canadian Immunization Research Network (CIRN) Provincial Collaborative Network (PCN) Investigators (2021). Effectiveness of BNT162b2 and mRNA-1273 covid-19 vaccines against symptomatic SARS-CoV-2 infection and severe covid-19 outcomes in Ontario, Canada: test negative design study. *BMJ 374*, n1943.

Davies, N. G., A. J. Kucharski, R. M. Eggo, A. Gimma, W. J. Edmunds, and Centre for the Mathematical Modelling of Infectious Diseases COVID-19 working group (2020). Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: a modelling study. *The Lancet Public Health 5*(7), e375–e385. Epub 2020-06-02.

Della Rossa, F., D. Salzano, A. Di Meglio, F. De Lellis, M. Coraggio, C. Calabrese, A. Guarino, R. Cardona-Rivera, P. De Lellis, D. Liuzza, F. Lo Iudice, G. Russo, and M. di Bernardo (2020). A network model of italy shows that intermittent regional strategies can alleviate the COVID-19 epidemic. *Nature Communications 11*, 5106.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2021). The macroeconomics of epidemics. *The Review of Financial Studies 34*(11), 5149–5187.

Fajgelbaum, P. D., A. Khandelwal, W. Kim, C. Mantovani, and E. Schaal (2021). Optimal lockdown in a commuting network. *American Economic Review: Insights 3*(4), 503–522.

Farboodi, M., G. Jarosch, and R. Shimer (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory 196*, 105293.

Fu, Z., G. Kaplan, B. Moll, and G. L. Violante (2020). The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the U.S. Working paper (author list as shown on authors' working-paper webpages), September 2020 version.

Fujii, D. and T. Nakata (2021). COVID-19 and output in Japan. *The Japanese Economic Review 72*(4), 609–650.

Fukao, M. and E. Shioji (2022). Is there a trade-off between COVID-19 control and economic activity? Implications from the Phillips Curve debate. *Asian Economic Policy Review 17*(1), 66–85.

Garriga, C., R. Manuelli, and S. Sanghi (2022). Optimal management of an epidemic: Lockdown, vaccine and value of life. *Journal of Economic Dynamics and Control 140*, 104351.

Glover, A., J. Heathcote, and D. Krueger (2022). Optimal age-based vaccination and economic mitigation policies for the second phase of the COVID-19 pandemic. *Journal of Economic Dynamics and Control 140*, 104306.

Goenka, A., L. Liu, and M.-H. Nguyen (2024). Modelling optimal lockdowns with waning immunity. *Economic Theory 77*, 197–234.

Hosono, K. (2021). Epidemic and economic consequences of voluntary and request-based lockdowns in Japan. *Journal of the Japanese and International Economies 61*, 101147.

Hritonenko, N. and Y. Yatsenko (2024). Analysis of optimal lockdown in integral economic–epidemic model. *Economic Theory 77*(1-2), 235–259.

Jones, C., T. Philippon, and V. Venkateswaran (2021). Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies 34*(11), 5188–5223.

Kamo, M., M. Murakami, and S. Imoto (2022). Effects of test timing and isolation length to reduce the risk of COVID-19 infection associated with airplane travel, as determined by infectious disease dynamics modeling. *Microbial Risk Analysis 20*, 100199.

Kissler, S. M., C. Tedijanto, E. Goldstein, Y. H. Grad, and M. Lipsitch (2020). Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. *Science 368*(6493), 860–868. Epub 2020-04-14.

Kubota, S. (2021). The macroeconomics of COVID-19 exit strategy: the case of Japan. *The Japanese Economic Review 72*(4), 651–682.

Loertscher, S. and E. V. Muir (2021). Road to recovery: Managing an epidemic. *Journal of Mathematical Economics 93*, 102482.

Ma, L., G. Shapira, D. de Walque, Q.-T. Do, J. Friedman, and A. A. Levchenko (2022). The intergenerational mortality trade-off of COVID-19 lockdown policies. *International Economic Review 63*(3), 1427–1468.

Miclo, L., D. Spiro, and J. Weibull (2022). Optimal epidemic suppression under an ICU constraint: An analytical solution. *Journal of Mathematical Economics 101*, 102669.

Moll, B. (2020). Lockdowns in SIR models. Lecture notes.

Morita, H. and T. Ono (2024). COVID-19 uncertainty index in Japan: Newspaper-based measures and economic activities. *International Review of Economics & Finance 93*, 390–403.

Pindyck, R. (2020). COVID-19 and the welfare effects of reducing contagion. Technical Report w27121, National Bureau of Economic Research, Cambridge, MA.

Prieur, F., W. Ruan, and B. Zou (2024). Optimal lockdown and vaccination policies to contain the spread of a mutating infectious disease. *Economic Theory 77*(1), 75–126.

Yang, H., Ö. Sürer, D. Duque, D. P. Morton, B. Singh, S. J. Fox, R. Pasco, K. Pierce, P. Rathouz, V. Valencia, Z. Du, M. Pignone, M. E. Escott, S. I. Adler, S. C. Johnston, and L. A. Meyers (2021). Design of COVID-19 staged alert systems to ensure healthcare capacity with minimal closures. *Nature Communications 12*, 3767.

Zardini, A., M. Galli, M. Tirani, D. Cereda, M. Manica, F. Trentini, G. Guzzetta, V. Marziano, R. Piccarreta, A. Melegaro, M. Ajelli, P. Poletti, and S. Merler (2021). A quantitative assessment of epidemiological parameters required to investigate COVID-19 burden. *Epidemics 37*, 100530.

Zhang, T. and H. Nishiura (2023). Estimating infection fatality risk and ascertainment bias of COVID-19 in Osaka, Japan from February 2020 to January 2022. *Scientific Reports 13*, 5540.