TCER Working Paper Series

Optimal Lockdown Policy with Virus Mutation

Quentin Batista
Masakazu Emoto
Naoki Maezono
Taisuke Nakata

March 2026

Working Paper E-221
https://www.tcer.or.jp/wp/pdf/e221.pdf

TOKYO CENTER FOR ECONOMIC RESEARCH
1-7-10-703 Iidabashi, Chiyoda-ku, Tokyo 102-0072, Japan

Abstract

We examine the implications of virus mutation for optimal lockdown policy in an epi-macro model. We consider three ways of modelling mutation---one deterministic setup and two stochastic setups featuring a two-state and three-state Markov process. We find that the effects of mutation on optimal lockdown policy are asymmetric. In particular, a future reduction in the transmission rate increases lockdown intensity by more than a future rise in the transmission rate lowers it. As a corollary to this asymmetry, an increase in uncertainty about future mutation is non-neutral and reduces lockdown intensity under the optimal policy.

Quentin Batista
Amazon Japan, Inc.
N/A
1-8-1 Shimomeguro, Meguro-ku, Tokyo 153-0064
batista.quent@gmail.com

Masakazu Emoto
Hitotsubashi University
Hitotsubashi Institute for Advanced Study
2-1 Naka, Kunitachi, Tokyo 186-8601
masakazu.emoto@gmail.com

Naoki Maezono
University of Tokyo
Graduate School of Public Policy
7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033
m7046ranpo464@g.ecc.u-tokyo.ac.jp

Taisuke Nakata
TCER
and
University of Tokyo
Faculty of Economics
7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033
taisuke.nakata@e.u-tokyo.ac.jp

Optimal Lockdown Policy with Virus Mutation*

Quentin Batista† — Amazon Japan
Masakazu Emoto‡ — Hitotsubashi University
Naoki Maezono§ — University of Tokyo
Taisuke Nakata¶ — University of Tokyo

This Draft: February 2026

Abstract

We examine the implications of virus mutation for optimal lockdown policy in an epi-macro model. We consider three ways of modelling mutation—one deterministic setup and two stochastic setups featuring a two-state and three-state Markov process. We find that the effects of mutation on optimal lockdown policy are asymmetric. In particular, a future reduction in the transmission rate increases lockdown intensity by more than a future rise in the transmission rate lowers it. As a corollary to this asymmetry, an increase in uncertainty about future mutation is non-neutral and reduces lockdown intensity under the optimal policy.

JEL: E32, E61, I18

Keywords: Epi-Macro Model, Lockdown, Macro-SIRD Model, Mutation, Optimal Policy, Uncertainty

*We thank Fumio Ohtake and seminar participants at the Center for Infectious Disease Education and Research, Osaka University, for their insightful comments. Taisuke Nakata is supported by JSPS Grant-in-Aid for Scientific Research (KAKENHI), Project Number 22H04927, and the Center for Advanced Research in Finance at the University of Tokyo. Quentin Batista worked on this project prior to his current employment.

†Amazon Japan, Inc.; Email: batista.quent@gmail.com.

‡Hitotsubashi Institute for Advanced Study, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo 186-8601; Email: masakazu.emoto@gmail.com.

§Graduate School of Public Policy, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033; Email: m7046ranpo464@g.ecc.u-tokyo.ac.jp.

¶Faculty of Economics and Graduate School of Public Policy, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033; Email: taisuke.nakata@e.u-tokyo.ac.jp.

1 Introduction

As with any infectious disease, SARS-CoV-2 has changed its viral characteristics over the course of the COVID-19 pandemic. From late 2020 to late 2021, several new variants emerged (alpha, delta, and omicron variants), each with distinct characteristics in terms of the transmission rate, severity rate, and fatality rate. These variants often caused new infection waves, forcing governments to adapt their policies to each variant's new viral characteristics.

In this paper, we examine the implications of virus mutation—a change in the transmission rate—for optimal lockdown policy. We conduct three distinct exercises. In the first exercise, we consider a model in which the path of the transmission rate is deterministic, but the rate is not constant: it rises or declines at a known future date. In the second exercise, we consider a model in which the transmission rate is stochastic and follows a two-state Markov process with an absorbing state. In the final exercise, we consider a model in which the transmission rate follows a three-state Markov process to examine the effect of mutation uncertainty. In the main text, we focus on mutation arising from a change in the transmission rate and point out several key takeaways. However, all the key takeaways are robust to two other parameters for viral characteristics—the fatality rate and the recovery rate—as discussed in the Appendix B.

In the first exercise using a deterministic model, we find that the effect of a future change in the transmission rate on the optimal lockdown policy is asymmetric. If the transmission rate is anticipated to rise in the future, it is optimal for the government to impose a less stringent lockdown than if the transmission rate is constant. By imposing a less stringent lockdown, the government could mitigate the decline in economic activities now and reduce the number of susceptible individuals before the transmission rate increases. If the transmission rate is anticipated to decline in the future, it is optimal for the government to impose a more stringent lockdown than if the transmission rate is constant. Thus, the effects of a decline or rise in the future transmission rate are qualitatively symmetric. However, we find that they are quantitatively asymmetric: that is, a decline in the future transmission rate increases lockdown intensity by more than the rise in the future transmission rate of the same magnitude. Intuitively, this asymmetry arises because of the non-negativity constraint on lockdown intensity: There is a limit to which a future decline in the transmission rate can reduce lockdown intensity because lockdown intensity cannot fall below zero.

In the second exercise using a model in which the transmission rate follows a two-state Markov process, we find a similar asymmetry in the effects of a possible future change in the transmission rate. As in the deterministic model, it is optimal for the government to impose a less (more) stringent lockdown if the transmission rate can rise (decline) in the future than if the transmission rate is constant. However, the nature of asymmetry somewhat differs from that in the deterministic model used for the first exercise.

In the third exercise using a model in which the transmission rate follows a three-state Markov process, we examine the effect of mutation uncertainty by comparing the optimal policy in this model with that in the model with a constant transmission rate. Reflecting the asymmetric effects of optimal lockdown intensity, we find that the effect of uncertainty is non-neutral. In particular, we find that it is optimal for the government to impose a less stringent lockdown in the presence of uncertainty.

Our paper builds on the literature analyzing optimal lockdown policy in epi-macro models. Though the analysis of epi-macro models predates the COVID-19 pandemic, this literature expanded during the COVID-19 crisis. Examples include Acemoglu et al. (2021), Alvarez et al. (2021), Berger et al. (2022), Chari et al. (2021), Eichenbaum et al. (2021), Farboodi et al. (2021), Fu et al. (2022), Glover et al. (2022), Gonzalez-Eiras and Niepelt (2025), Piguillem and Shi (2022), Rachel (2026), among many others. This literature has largely focused on a deterministic environment with constant parameter values, except for the possibility of vaccine or treatment arrival that would practically end the pandemic. Our contribution is to examine the implications for the optimal policy of mutation that would discretely change parameters governing the virus characteristics.

Our first and second analyses build on existing analyses mentioned above that examine how the expectation of vaccine or treatment arrival in the future affects current lockdown policy. A special case of our analyses in which the transmission rate declines to zero (or the recovery rate rises to infinity) at some point in the future would correspond to the analysis in the literature in which vaccines (or treatment) arrive at some point in the future and their efficacy is perfect. See Acemoglu et al. (2020), Assenza et al. (2021), Bodenstein et al. (2022), Boppart et al. (2025), Eichenbaum et al. (2021), Garriga et al. (2022), Glover et al. (2023), Jones et al. (2021), Makris and Toxvaerd (2020), among others, for this type of analysis. The marginal contribution of our first and second analyses relative to these analyses is that we point out the asymmetry in the effect of expected changes in the parameters governing virus characteristics, which leads to non-neutrality of uncertainty found in our third exercise.[1]

Our third analysis on uncertainty based on a three-state Markov process builds on several papers that explicitly consider either time-variation or uncertainty in the parameters characterizing the virus. Barnett et al. (2023) allow for uncertainty in various parameters of an epi-macro model and characterize max-min policy—robust policy in the language of

[1]Though not focus of our paper, our first and second analyses show that the effect of a future change in the rate depends on whether you model it in a deterministic way or in a stochastic way using a two-state Markov process. This result can be seen as another marginal contribution of our analysis relative to existing studies.

Hansen and Sargent—in it, finding that the government imposes more strict lockdowns under parameter uncertainty. Our analysis differs from their work because we assume that the government maximizes the expected value of welfare, as opposed to maximizing welfare under the worst-case scenario, and consider a different shock structure. Bandyopadhyay et al. (2021) study optimal lockdown policy in a model in which the government is uncertain about the transmission rate and early lockdown has the benefit of promoting good habit formation among citizens for social distancing, but has the cost of not being able to learn about the transmission rate. Our paper differs from this paper in that we examine the implication of a specific type of parameter uncertainty—mutation uncertainty—in a standard framework abstracting from habit formation and learning.

Federico and Ferrari (2021) study optimal lockdown policy in an epi-macro model in which the transmission rate follows a stochastic diffusive process, finding that an increase in the fluctuations of the transmission rate gives rise to a lockdown that begins earlier but is diluted over a longer period of time. We consider an alternative shock structure intended to capture uncertainty arising from mutation that occurs less frequently. Gollier (2020) studies optimal lockdown policy in an exponential-decay model of infectious diseases and shows that uncertainty about the reproduction number reduces the optimal initial rate of lockdown. We differ from this paper because we examine uncertainty in a more standard SIRD model.

Finally, Prieur et al. (2024) study optimal lockdown and vaccination policy in an SIS model in which there is uncertainty in both the timing and the size of a change in the transmission rate, finding that uncertainty surrounding future mutation of the disease expedites lockdown intervention whenever mutation increases contagiousness. Their work is closest to ours in that they consider uncertainty arising from mutation. Our work differs from theirs because we adopt an SIRD model, instead of an SIS model, and consider both two-state and three-state Markov processes.

The rest of the paper is organized as follows. Section 2 presents the model, formulates the government's optimization problem, and discusses parameter values. Section 3 presents the results. Section 4 concludes.

2 Model

2.1 SIRD Model with Production

We use an infinite-horizon SIRD model formulated in discrete time. In the deterministic version of the model, we will drop the expectations operator from the model.

The dynamics of the SIRD model are given by the following equations:

$$E_t S_{t+1} = S_t - N_t \tag{1}$$
$$E_t I_{t+1} = I_t + N_t - N_t^{IR} - N_t^{ID} \tag{2}$$
$$E_t R_{t+1} = R_t + N_t^{IR} \tag{3}$$
$$E_t D_{t+1} = D_t + N_t^{ID} \tag{4}$$
$$N_t^{IR} = \gamma_t I_t \tag{5}$$
$$N_t^{ID} = \delta_t I_t \tag{6}$$
$$POP = S_t + I_t + R_t + D_t. \tag{7}$$

S_t, I_t, R_t, D_t denote the number of susceptible, infected, recovered, and deceased, respectively. The flow variables N_t, N_t^{IR}, and N_t^{ID} are the number of newly infected, newly recovered, and newly deceased between time t and time $t+1$, respectively. Parameters γ_t and δ_t denote recovery rate and fatality rate, respectively. The total population is a constant denoted by POP. In the main body of the paper, we assume that γ_t and δ_t are constant over time, which we denote by γ and δ. In the Appendix, we will consider models in which either γ_t or δ_t is time-varying.

The matching function for new infections follows:

$$N_t = \frac{1}{POP} \tilde{\beta}_t I_t S_t \tag{8}$$

where

$$\tilde{\beta}_t = \beta_t (1 - h\alpha_t)^2, \tag{9}$$

and α_t denotes lockdown intensity. β_t and $\tilde{\beta}_t$ denote the transmission rate and the lockdown-adjusted transmission rate, respectively. h denotes lockdown effectiveness.

We assume that the output of the economy y_t is determined by the following production function:

$$y_t = (1 - \alpha_t)(S_t + \omega_I I_t + R_t) \tag{10}$$

where ω_I denotes the productivity of infected workers.

2.2 Government's Optimization Problem

The government's optimization problem is to choose a state-contingent sequence of lockdown intensities $\{\alpha_t(\beta^t)\}_{t=0}^{\infty}$ to maximize the social welfare:

$$\max_{\{\alpha_t(\beta^t)\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \rho^t \left(y_t(\beta^t) - \chi N_t^{ID}(\beta^t) \right), \tag{11}$$

subject to equations (1)–(10) and $\alpha_t \in [0,1]$.[2] χ converts deaths into output-equivalent welfare losses (i.e., the value of statistical life in model units). ρ is the discount factor of the government. Following the standard notation in modern macroeconomics, we use x^t to denote the history of a random variable x from time 0 to time t. When there is uncertainty, the optimal policy and allocations under it are state-contingent sequences. In the objective function above, we made the state-contingent nature of the solution explicit by explicitly stating that each variable at time t is a function of the history of shocks up to time t.

2.3 Shock Processes

In all three exercises, the baseline shock process is a constant path. We compute the optimal policy with mutation in three different ways, and contrast it with that under the constant-path case.

In the first exercise, the transmission rate under mutation is deterministic, but changes its value at a fixed date T. The change is anticipated by the government at the beginning of time 0. In the second exercise, the transmission rate follows a two-state Markov process with an absorbing state. The transition matrix is given by:

$$\Pi_2 = \begin{bmatrix} 1-p_2 & p_2 \\ 0 & 1 \end{bmatrix} \tag{12}$$

where p_2 denotes the probability of mutation for a two-state Markov process.

In the third exercise, the transmission rate follows a three-state Markov process with absorbing states. The three-state Markov process is symmetric. The probability of a rise in the transmission rate is the same as that of a decline. The size of the rise is the same as

[2] See Appendix A for details on the numerical methods used to solve the government's optimization problem.

that of the decline. The transition matrix is given by:

$$\Pi_3 = \begin{bmatrix} 1 & 0 & 0 \\ p_3 & 1-2p_3 & p_3 \\ 0 & 0 & 1 \end{bmatrix}. \tag{13}$$

where p_3 denotes the probability of mutation for a three-state Markov process.

In all three exercises, we consider both a rise and a decline in the transmission rate.

2.4 Parameter Values

We interpret a period in our model as a week. We set $\gamma = 7/18$. Together with $\delta = 0.01$, this implies an average infectious duration of $1/(\gamma + \delta)$ weeks (about 18 days), consistent with Eichenbaum et al. (2021). We set the fatality rate $\delta = 0.01$. This value is higher than the population-averaged infection fatality risk associated with COVID-19, as reported in Zhang and Nishiura (2023). We choose this higher number for the ease of exposition. All of the key takeaways in our paper are robust to alternative fatality rates. The base transmission rate (β) is set to one, so that the implied basic reproduction number is 2.5, consistent with the Centers for Disease Control and Prevention (2021). We set the discount factor $\rho = 0.999$, following Eichenbaum et al. (2021). We set χ to 1000 so that the implied value of a statistical life in our model is consistent with the calibration target in Farboodi et al. (2021). We set the probability of mutation as $p_2 = 0.025$ or $p_3 = 0.0125$, so the expected duration until mutation is 40 weeks.

Table 1: Parameter Values

Parameter	Description	Value
POP	Total population size	100
γ	Recovery rate	7/18
δ	fatality rate	0.01
β	Base transmission rate	1
ρ	Discount factor for future costs	0.999
χ	Cost weight for deaths relative to economic cost	1,000
h	Lockdown effectiveness	1
ω_I	productivity of infected workers.	0.8
p_2	The probability of mutation in the two-state Markov process	0.025
p_3	The probability of mutation in the three-state Markov process	0.0125

Note: See Appendix C for robustness check regarding the productivity of infected workers, ω_I.

3 Results

3.1 An Anticipated Change at a Fixed Date

In this exercise, we examine the implications of an anticipated future change in the transmission rate for optimal lockdown policy. We do so by comparing (i) the economy with a constant transmission rate (the economy without mutation) and (ii) the economy in which the transmission rate changes its value at a fixed period T (the economy with anticipated mutation). We consider both an increase and a decrease in the transmission rate.

Figure 1 shows the results. In the figure, solid black lines show the economy without mutation, whereas solid red and blue lines show the economy with an anticipated increase and decrease in the transmission rate at time T, respectively.

Figure 1: Dynamics with and without an Anticipated Change at $T = 40$: Transmission Rate (β)

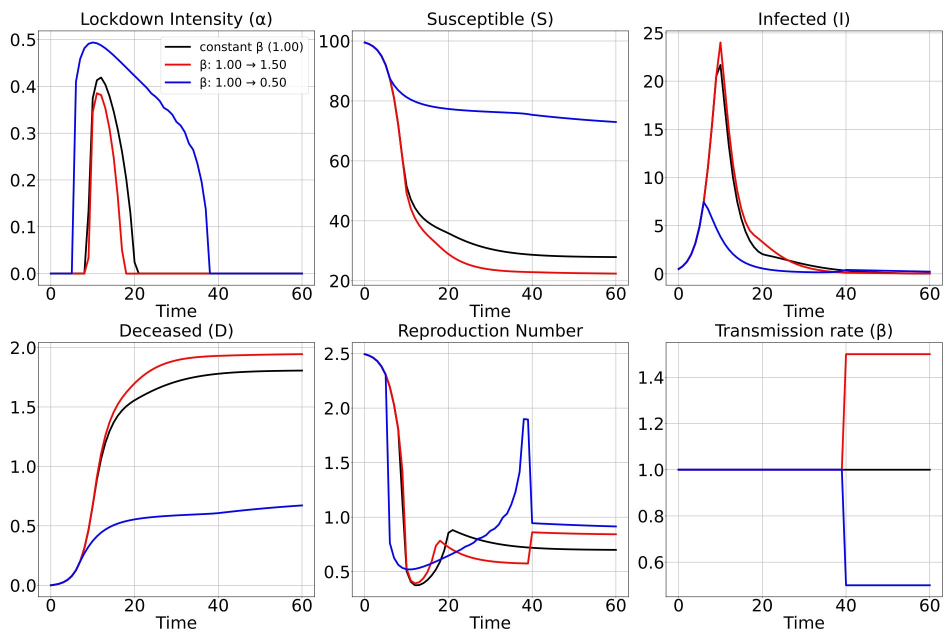

Note: The black line is the constant-β benchmark with $\beta = 1.00$. The red and blue lines are realizations under an anticipated change in β, where agents expect β to remain at 1.00 until $T = 40$ and then permanently switch to the high value ($\beta = 1.50$, red) or the low value ($\beta = 0.50$, blue). All simulations start from the same initial condition and differ only in the anticipated change in β at $T = 40$.

In the economy without mutation, the government imposes lockdown from time 9 to time 20 to mitigate the infection wave, as shown by the solid black line in the top left panel. Comparing the solid black lines with the solid red lines, we see that the anticipation of a future increase in the transmission rate leads to a less stringent lockdown. If a second infection wave occurs after the transmission rate increases, the cost of managing the pandemic becomes high. That is, the government would need to impose a strict lockdown to contain the wave. To reduce the possibility of the second infection wave, or the size of it if it occurs, the government has an incentive to reduce the pool of susceptible individuals at the time the transmission rate increases. To achieve this, the government aims to reduce the severity of the lockdown during the first wave of infection so that more people become infected before the transmission rate rises.

Comparing the solid black with blue lines, we see that the anticipation of a future decline in the transmission rate leads to a more stringent lockdown. The logic is the opposite of the other case. Once the transmission rate declines, the cost of managing the pandemic becomes low. That is, the government would not need to impose a strict lockdown to control the wave. Thus, the government has less incentive to reduce the pool of susceptible persons before the transmission rate declines. To achieve that, the government would like to impose a stricter lockdown. This result is consistent with the findings from the existing studies that have studied the implication of the promise of vaccines on the optimal lockdown policy, such as Acemoglu et al. (2020), Assenza et al. (2021), Boppart et al. (2025), Eichenbaum et al. (2021), Garriga et al. (2022), and Glover et al. (2023).

From the figure, we see that the effects of an anticipated increase and decrease in the transmission rate are asymmetric. In particular, an anticipated decrease in the transmission rate increases the lockdown intensity by more than an anticipated increase in the transmission rate reduces it. Intuitively, this asymmetry arises because of the non-negativity constraint on lockdown intensity: There is a limit to which a future increase in the transmission rate can reduce lockdown intensity because lockdown intensity cannot fall below zero.

To further highlight the asymmetry, we plot how the path of lockdown intensity varies with alternative shock sizes. In Figure 2, we show three alternative shock sizes. The left (right) panel shows the path of optimal lockdown intensity when the shock size is smaller (larger) than in the baseline. The middle panel is the same as in the baseline. As the absolute size of the anticipated *increase* in the transmission rate becomes larger, the optimal lockdown intensity declines in a relatively linear way. However, as the absolute size of the anticipated *decline* in the transmission rate becomes larger, the optimal lockdown intensity increases, but increases in a nonlinear manner.

To quantify the degree of asymmetry, Table 2 shows the cumulative lockdown intensity for all cases considered in Figure 2. As the terminal β declines from the baseline value of one, the cumulative lockdown intensity increases steeply. On the other hand, as the terminal β increases from the baseline value of one, the cumulative lockdown intensity declines, but by a smaller amount.

Figure 2: Optimal Lockdown Policy under an Anticipated Change at $T = 40$:
Transmission Rate (β)—Alternative Shock Sizes—

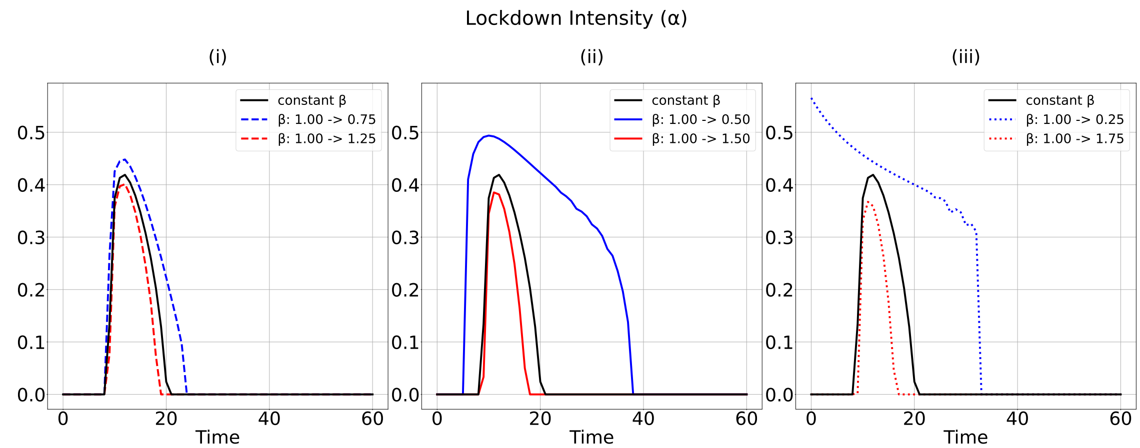

Note: This figure plots the simulated path of lockdown intensity (α) under a constant transmission rate β and anticipated, permanent changes in β occurring at $T = 40$. In the benchmark (black solid line), $\beta = 1.00$ for all t. In the alternative scenarios, agents anticipate that β remains at 1.00 until $T = 40$ and then permanently switches to a new level, as indicated in each panel's legend. Panel (i) compares $\beta : 1.00 \rightarrow 0.75$ (blue dashed) and $\beta : 1.00 \rightarrow 1.25$ (red dashed); panel (ii) compares $\beta : 1.00 \rightarrow 0.50$ (blue solid) and $\beta : 1.00 \rightarrow 1.50$ (red solid); panel (iii) compares $\beta : 1.00 \rightarrow 0.25$ (blue dotted) and $\beta : 1.00 \rightarrow 1.75$ (red dotted). All simulations start from the same initial condition and differ only in the anticipated path of β.

Table 2: Cumulative Lockdown Intensity under an Anticipated Change at $T = 40$:
Transmission Rate (β)

$\beta : 1.00 \rightarrow 0.25$	$\beta : 1.00 \rightarrow 0.50$	$\beta : 1.00 \rightarrow 0.75$	constant β	$\beta : 1.00 \rightarrow 1.25$	$\beta : 1.00 \rightarrow 1.50$	$\beta : 1.00 \rightarrow 1.75$
14.1	12.4	4.7	3.4	2.8	2.3	1.9

Note: This table reports cumulative lockdown intensity, defined as $\sum_t \alpha_t$ over the sample period, under a constant transmission rate β and six scenarios with an anticipated change in β at $T = 40$. In the benchmark, $\beta = 1.00$ throughout. In the anticipated-change cases, β starts from 1.00 and, in the realization summarized here, changes at $T = 40$ from 1.00 to 0.25, 0.50, 0.75, 1.25, 1.50, or 1.75 as indicated by the column headings; after the change, β remains at the new level. All simulations start from the same initial condition and differ only in the assumed path of β.

3.2 A Two-State Markov Process

In this exercise, we examine the implications of a possible future change in the transmission rate for optimal lockdown policy. We do so by comparing (i) the economy with a constant transmission rate (the economy without mutation) and (ii) the economy in which the transmission rate follows a two-state Markov process with an absorbing state (the economy with possible mutation).

Figure 3 shows the results. In the figure, solid black lines are for the economy without mutation, whereas solid red and blue lines are for the economy with possible mutation—a rise and a decline in the transmission rate—respectively.

Figure 3: Dynamics with Constant-Parameter vs. Two-State Markov Process: Transmission Rate (β)

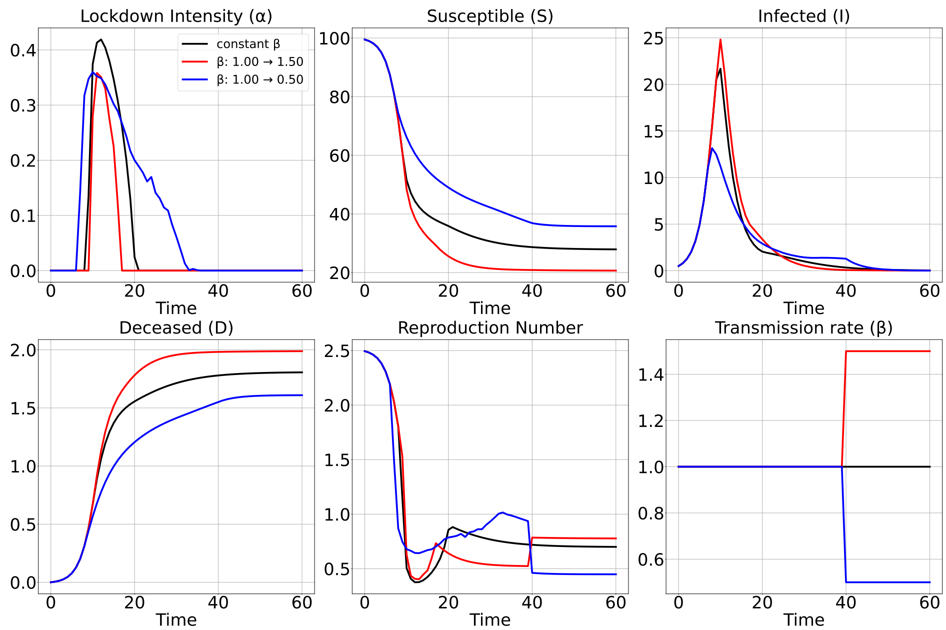

Note: The black line is the benchmark with constant $\beta = 1.00$. The red and blue lines are realizations under a two-state Markov process for β in which β starts from 1.00 and, from $T = 40$, can switch to an alternative (absorbing) state with per-period transition probability $p_2 = 0.025$. In the realizations shown, the switch occurs at $T = 40$, moving to the high state ($\beta = 1.50$, red) or the low state ($\beta = 0.50$, blue); after switching, β remains at the new level.

Let us first discuss the economy with a possible increase in the transmission rate, as shown by the red lines. As in the deterministic model with an increase in the transmission rate at a fixed date, it is optimal for the government to impose less stringent lockdowns if the transmission rate can increase in the future than otherwise.

Let us now discuss the economy with a possible decline in the transmission rate, as shown by the blue line. As in the deterministic model with a decrease in the transmission rate at a fixed date, it is optimal for the government to impose more stringent lockdowns if the transmission rate can decline in the future than otherwise. However, the optimal lockdown path is not uniformly higher than in the model without mutation over time. For some time periods around the peak of the lockdown intensity in the model with a constant transmission rate, the lockdown intensity is lower in the model with a possible decline in the transmission rate than in the model without mutation.

To see the asymmetry from another perspective, we plot how the path of lockdown intensity varies with the size of unanticipated change. In Figure 4, we show the path of optimal lockdown intensity under two different-sized shocks—one smaller and one larger—in the left and right panels, respectively. The middle panel is for the baseline shock size. As the absolute size of the anticipated increase in the transmission rate gets larger, the optimal lockdown intensity declines in a relatively linear way. However, as the absolute size of the anticipated decline in the transmission rate gets larger, the optimal lockdown intensity increases in a nonlinear way. To quantify the degree of asymmetry, Table 3 shows the cumulative lockdown intensity under these different shock sizes, as in Table 2.

The paths for future possible increases in the transmission rate in Figure 2 are qualitatively similar to those for the anticipated increase in the transmission rate in Figure 1. Interestingly, the paths for future possible declines in the transmission rate in Figure 2 are qualitatively different from those for the anticipated decline in the transmission rate in Figure 1. That is, the nature of asymmetry depends on whether we model mutation via an anticipated change or a two-state Markov process.

Figure 4: Optimal Lockdown Policy under a Two-State Markov Process: Transmission Rate (β)

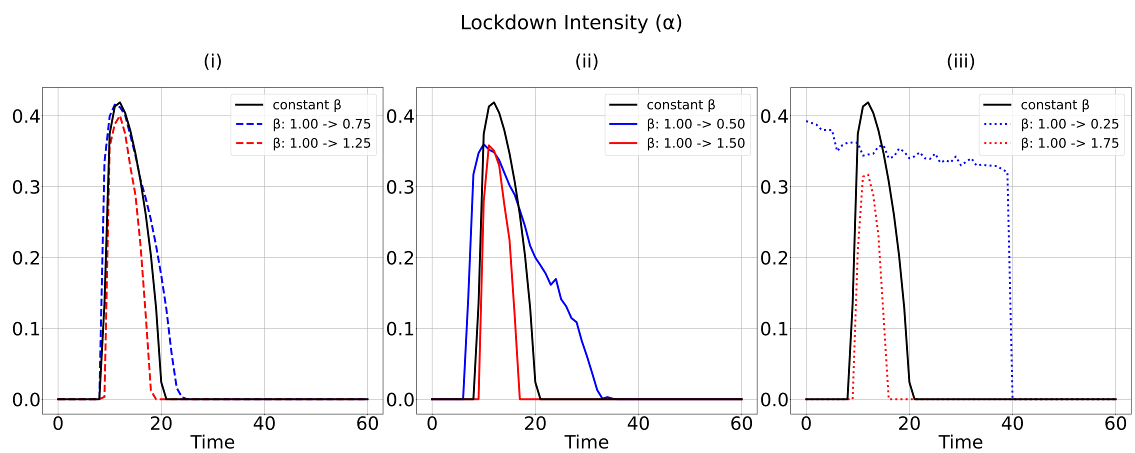

Note: This figure plots the simulated path of lockdown intensity α_t under a constant transmission rate β and under a two-state Markov process for β. In the benchmark (black solid line), $\beta = 1.00$ throughout. In the two-state Markov cases, β starts at 1.00 and, from $t = 0$, can switch to an alternative (absorbing) state with per-period transition probability $p_2 = 0.025$; after switching, β remains at the new level. The legend reports the alternative-state value. In the realizations shown, the switch occurs at $T = 40$. Panel (i) compares $\beta : 1.00 \rightarrow 0.75$ (blue dashed) and $\beta : 1.00 \rightarrow 1.25$ (red dashed); panel (ii) compares $\beta : 1.00 \rightarrow 0.50$ (blue solid) and $\beta : 1.00 \rightarrow 1.50$ (red solid); panel (iii) compares $\beta : 1.00 \rightarrow 0.25$ (blue dotted) and $\beta : 1.00 \rightarrow 1.75$ (red dotted). All simulations start from the same initial condition and differ only in the assumed process for β.

Table 3: Cumulative Lockdown Intensity under a Two-State Markov Process: Transmission Rate (β)

$\beta: 1.00 \rightarrow 0.25$	$\beta: 1.00 \rightarrow 0.50$	$\beta: 1.00 \rightarrow 0.75$	constant β	$\beta: 1.00 \rightarrow 1.25$	$\beta: 1.00 \rightarrow 1.50$	$\beta: 1.00 \rightarrow 1.75$
13.9	5.4	4.1	3.4	2.5	1.9	1.5

Note: This table reports cumulative lockdown intensity, defined as $\sum_t \alpha_t$ over the sample period, under a constant transmission rate β and six two-state Markov cases for β. In the benchmark constant-parameter case, $\beta = 1.00$. In the two-state Markov cases, β starts from 1.00 and, from $t = 0$, can switch to an alternative (absorbing) state with per-period transition probability $p_2 = 0.025$. The alternative-state value is 0.25, 0.50, 0.75, 1.25, 1.50, or 1.75, as indicated by the column headings. In the realizations summarized here, the switch occurs at $T = 40$; after switching, β remains at the new level. All simulations start from the same initial condition and differ only in the assumed process for β.

3.3 A Three-State Markov Process

In this exercise, we examine the implications of mutation uncertainty for optimal lockdown policy. We do so by comparing (i) the economy with a constant transmission rate (the economy without mutation uncertainty) and (ii) the economy in which the transmission rate follows a three-state Markov process with absorbing states (the economy with mutation uncertainty).

Figure 5 shows the results. In the figure, solid black and red lines are for the economy without and with uncertainty. According to the figure, uncertainty is not neutral. In particular, it is optimal for the government to impose a less stringent lockdown in the presence of mutation uncertainty than in its absence. Non-neutrality of uncertainty—or breakdown of certainty equivalence—makes sense in light of the asymmetry discussed in the previous two models.

Figure 5: The Effects of Mutation Uncertainty on Optimal Lockdown Policy: Transmission Rate (β)

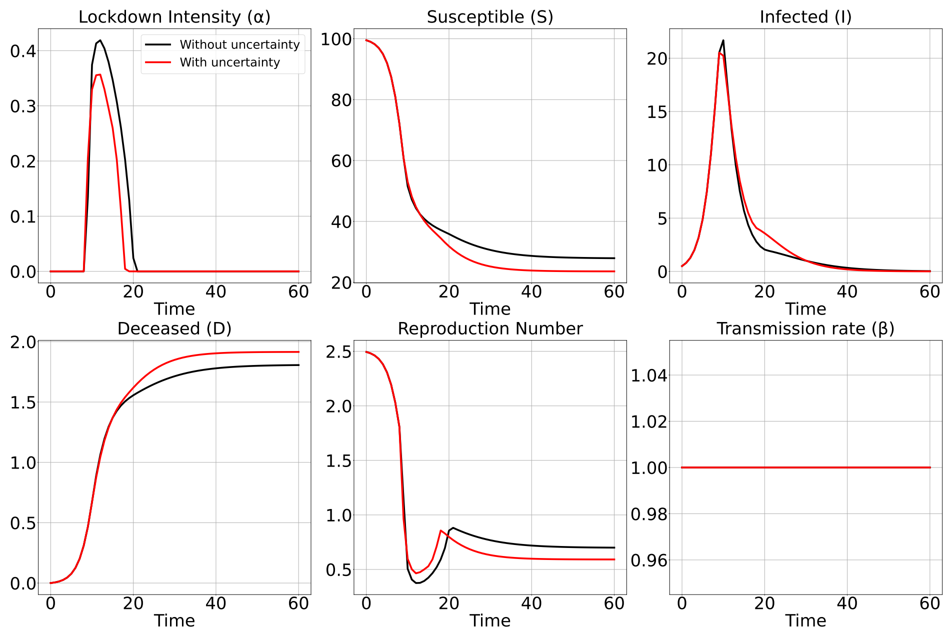

Note: The black line (without uncertainty) shows the benchmark case with a constant transmission rate at $\beta = 1.00$. The red line (with uncertainty) shows a specific realization of the three-state Markov model in which β starts from 1.00 and, from $t = 0$, can switch in any period to the high state ($\beta = 1.50$) or the low state ($\beta = 0.50$) with per-period transition probability $p_3 = 0.0125$. In the realization shown, $\beta = 1$ over the entire horizon. All simulations start from the same initial condition and differ only in the assumed process for β.

Why does the government reduce—as opposed to increase—lockdown intensity in the presence of uncertainty? To answer this question, it is useful to consider two types of mistakes the government could face in this model with a three-state Markov shock. The first type of mistake is that the government intensifies lockdown, focusing on the possible future decline in the transmission rate, but the transmission rate ends up increasing in the future. The second type of mistake is the opposite: the government relaxes lockdown, focusing on the possible future increase in the transmission rate, but the transmission rate ends up declining in the future.

The first mistake turns out to be more costly than the second mistake. In the first mistake, when the transmission rate increases at some point in the future, the government has already incurred economic costs that are substantially larger than in the constant-parameter case. Furthermore, once the transmission rate increases, there will be a large number of deaths from that point on because there are many susceptible people left due to a stringent lockdown up to that point. In the second mistake, when the transmission rate declines at some point in the future, the government has allowed infections and infection-induced deaths to rise more than in the constant-parameter case. Once the transmission rate declines, there will not be many additional deaths from that point on because the number of susceptible people has already declined by a large amount. All told, the first mistake is associated with substantially higher economic costs without necessarily reaping much gain in terms of cumulative deaths than in the constant-parameter case, whereas the second mistake is associated with some reduction in economic costs without necessarily implying substantially higher costs in terms of cumulative deaths than in the constant-parameter case.

4 Conclusion

In this paper, we examined the implications of virus mutation for optimal lockdown policy in an epi-macro model. We considered three ways of modelling mutation—a deterministic setup and two stochastic setups featuring a two-state or three-state Markov process. We highlighted the following takeaways. First, the effects of mutation are asymmetric. In particular, in both the model with an anticipated shock at a fixed date and the model with a two-state Markov shock, we saw that a future reduction in the transmission rate increases lockdown intensity by more than a future rise in the transmission rate lowers it. Second, as a corollary to this asymmetry, an increase in uncertainty about future mutation is non-neutral and reduces lockdown intensity under the optimal policy.

We focused on a particular type of uncertainty in this paper. In reality, the government faces a myriad of uncertainties in a pandemic. The government may be uncertain about data, the appropriate model, and its citizens' priority over lives versus livelihoods. As discussed in the introduction, some researchers have analyzed the implications of uncertainty for the conduct of lockdown policy, but such research is still limited in its number. Further research on this issue is likely to provide invaluable insights to future policymakers in the next pandemic.

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2020). A multi-risk sir model with optimally targeted lockdown. NBER Working Paper Series 27102, National Bureau of Economic Research.

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 lockdown, testing, and tracing. *American Economic Review: Insights 3*(3), 367–82.

Assenza, T., C. Hellwig, F. Collard, M. Dupaigne, P. Fève, S. Kankanamge, and N. Werquin (2021). The hammer and the dance: Equilibrium and optimal policy during a pandemic crisis. Working Papers hal-03186935, HAL open science.

Bandyopadhyay, S., K. Chatterjee, K. Das, and J. Roy (2021). Learning versus habit formation: Optimal timing of lockdown for disease containment. *Journal of Mathematical Economics 93*, 102452.

Barnett, M., G. Buchak, and C. Yannelis (2023). Epidemic responses under uncertainty. *PNAS 120*(2), e2208111120.

Berger, D., K. Herkenhoff, C. Huang, and S. Mongey (2022). Testing and reopening in an seir model. *Review of Economic Dynamics 43*, 1–21.

Bodenstein, M., G. Corsetti, and L. Guerrieri (2022). Social distancing and supply disruptions in a pandemic. *Quantitative Economics 13*(2), 681–721.

Boppart, T., K. Harmenberg, J. Hassler, P. Krusell, and J. Olsson (2025). Integrated epi-econ assessment: Quantitative theory. *Quantitative Economics 16*(1), 89–131.

Centers for Disease Control and Prevention (2021, March). Covid-19 pandemic planning scenarios.

Chari, V., R. Kirpalani, and C. Phelan (2021). The hammer and the scalpel: On the economics of indiscriminate versus targeted isolation policies during pandemics. *Review of Economic Dynamics 42*, 1–14.

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

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

Federico, S. and G. Ferrari (2021). Taming the spread of an epidemic by lockdown policies. *Journal of Mathematical Economics 93*, 102453.

Fu, Y., H. Jin, H. Xiang, and N. Wang (2022). Optimal lockdown policy for vaccination during covid-19 pandemic. *Finance Research Letters 45*, 102123.

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.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Ríos-Rull (2023). Health versus wealth: On the distributional effects of controlling a pandemic. *Journal of Monetary Economics 140*, 34–59.

Gollier, C. (2020). Pandemic economics: optimal dynamic confinement under uncertainty and learning. *Geneva Risk and Insurance Review 45*, 80–93.

Gonzalez-Eiras, M. and D. Niepelt (2025). A tractable model of epidemic control and equilibrium dynamics. *Journal of Economic Dynamics and Control 178*, 105145.

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.

Makris, M. and F. Toxvaerd (2020). Great expectations: Social distancing in anticipation of pharmaceutical innovations. Cambridge Working Papers in Economics 2097, Faculty of Economics, University of Cambridge.

Piguillem, F. and L. Shi (2022). Optimal covid-19 quarantine and testing policies. *The Economic Journal 132*(647), 2534–2562.

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*, 75–126.

Rachel, L. (2026). An analytical model of behavior and policy in an epidemic. *American Economic Journal: Macroeconomics*. Forthcoming.

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.

Online Appendices

A Solution Methods

A.1 Model with an Anticipated Parameter-Change

We employ a brute-force (direct) optimization approach that treats the period-by-period lockdown intensity sequence as the object of choice. The method fixes a finite horizon and optimizes the lockdown intensity in each period jointly, rather than computing a state-contingent policy rule over the full state space.

The algorithm proceeds in two main steps: (i) forward simulation, where for a given candidate period-by-period lockdown intensity sequence we simulate the SIRD dynamics period by period starting from an initial condition and obtain the implied trajectory of health outcomes, and (ii) sequence update, where we minimize the discounted objective over the period-by-period lockdown intensity sequence using a gradient method with automatic differentiation subject to feasibility. After each update, we enforce the lockdown bounds by projecting each period's lockdown intensity back to the admissible interval.

A.2 Model with a Two- and Three-State Markov Process

We solve the stochastic model by policy function iteration on a discretized state space $x = (\theta, I, R + D)$, where $\theta \in \{\beta, \gamma, \delta\}$. The control is the lockdown intensity $\alpha \in [0, 1]$. Given a candidate policy $\pi(x)$, we perform policy evaluation by iterating the policy-specific Bellman operator until the value function V^π converges. The continuation value is computed by taking expectations over θ' using the Markov transition matrix and interpolating V^π on the $(\theta, I, R + D)$ grid at the implied next state. We then conduct policy improvement by maximizing the Bellman objective over a discretized grid for α at each state node, updating π. We iterate evaluation and improvement until convergence.

B Recovery and Fatality Rates

In the main body of the paper, we focused on the transmission rate. All the key takeaways are robust to allowing mutation that leads to a change in the recovery and fatality rate.

Table 4: Cumulative Lockdown Intensity under an Anticipated Change at $T = 40$
Recovery Rate (γ)

$\gamma: 7/18 \to 2/18$	$\gamma: 7/18 \to 4/18$	$\gamma: 7/18 \to 6/18$	constant γ	$\gamma: 7/18 \to 8/18$	$\gamma: 7/18 \to 10/18$	$\gamma: 7/18 \to 12/18$
0.5	1.8	2.9	3.4	4.0	6.6	10.3

Note: This table reports cumulative lockdown intensity, defined as $\sum_t \alpha_t$ over the sample period, under a constant recovery rate γ and six scenarios with an anticipated change in γ at $T = 40$. In the benchmark, $\gamma = 7/18$ throughout. In the anticipated-change cases, γ starts from 7/18 and, in the realization summarized here, changes at $T = 40$ from 7/18 to 2/18, 4/18, 6/18, 8/18, 10/18, or 12/18 as indicated by the column headings; after the change, γ remains at the new level. All simulations start from the same initial condition and differ only in the assumed path of γ.

Table 5: Cumulative Lockdown Intensity under a Two-State Markov Process:
Recovery Rate (γ)

$\gamma: 7/18 \to 2/18$	$\gamma: 7/18 \to 4/18$	$\gamma: 7/18 \to 6/18$	constant γ	$\gamma: 7/18 \to 8/18$	$\gamma: 7/18 \to 10/18$	$\gamma: 7/18 \to 12/18$
0.0	1.2	2.7	3.4	3.7	4.4	5.0

Note: This table reports cumulative lockdown intensity, defined as $\sum_t \alpha_t$ over the sample period, under a constant recovery rate γ and six two-state Markov cases for γ. In the benchmark, $\gamma = 7/18$ throughout. In the two-state Markov cases, γ starts from 7/18 and, from $t = 0$, can switch to an alternative (absorbing) state with per-period transition probability $p_2 = 0.025$. The alternative-state value is 2/18, 4/18, 6/18, 8/18, 10/18, or 12/18, as indicated by the column headings. In the realizations summarized here, the switch occurs at $T = 40$; after switching, γ remains at the new level. All simulations start from the same initial condition and differ only in the assumed process for γ.

Table 6: Cumulative Lockdown Intensity under an Anticipated Change at $T = 40$
Fatality Rate (δ)

$\delta: 0.010 \to 0.003$	$\delta: 0.010 \to 0.005$	$\delta: 0.010 \to 0.007$	constant δ	$\delta: 0.010 \to 0.013$	$\delta: 0.010 \to 0.015$	$\delta: 0.010 \to 0.017$
12.47	4.00	3.68	3.40	3.21	3.11	3.03

Note: This table reports cumulative lockdown intensity, defined as $\sum_t \alpha_t$ over the sample period, under a constant fatality rate δ and six scenarios with an anticipated change in δ at $T = 40$. In the benchmark, $\delta = 0.01$ throughout. In the anticipated-change cases, δ starts from 0.01 and, in the realization summarized here, changes at $T = 40$ from 0.01 to 0.003, 0.005, 0.007, 0.013, 0.015, or 0.017 as indicated by the column headings; after the change, δ remains at the new level. All simulations start from the same initial condition and differ only in the assumed path of δ.

Table 7: Cumulative Lockdown Intensity under a Two-State Markov Process: Fatality Rate (δ)

$\delta: 0.010 \to 0.003$	$\delta: 0.010 \to 0.005$	$\delta: 0.010 \to 0.007$	constant δ	$\delta: 0.010 \to 0.013$	$\delta: 0.010 \to 0.015$	$\delta: 0.010 \to 0.017$
3.97	3.58	3.42	3.40	3.19	3.18	3.18

Note: This table reports cumulative lockdown intensity, defined as $\sum_t \alpha_t$ over the sample period, under a constant fatality rate δ and six two-state Markov cases for δ. In the benchmark, $\delta = 0.01$ throughout. In the two-state Markov cases, δ starts from 0.01 and, from $t = 0$, can switch to an alternative (absorbing) state with per-period transition probability $p_2 = 0.025$. The alternative-state value is 0.003, 0.005, 0.007, 0.013, 0.015, or 0.017, as indicated by the column headings. In the realizations summarized here, the switch occurs at $T = 40$; after switching, δ remains at the new level. All simulations start from the same initial condition and differ only in the assumed process for δ.

C Robustness Analyses

This Appendix examines the sensitivity of our main results to the assumption regarding the relative productivity of infected individuals, ω_I. In our benchmark analysis, we set $\omega_I = 0.8$, assuming that those infected are slightly less productive than susceptible or recovered individuals.

The primary takeaway from Tables 8 and 9 is that the optimal lockdown policy is remarkably robust to variations in this parameter. Whether we assume infected individuals lose all productivity $\omega_I = 0.0$ or remain as productive as healthy individuals $\omega_I = 1.0$, the cumulative lockdown intensity $\sum \alpha_t$ remains nearly identical across all scenarios. This suggests that the government's decision is primarily driven by the transmission dynamics β_t and the associated health costs, rather than the marginal loss in labor productivity of the currently infected.

Table 8: Cumulative Lockdown Intensity under an Anticipated change at $T = 40$: Transmission Rate (β)—Alternative Values for ω_I—

ω_I	$\beta: 1.00 \to 0.25$	$\beta: 1.00 \to 0.50$	$\beta: 1.00 \to 0.75$	constant β	$\beta: 1.00 \to 1.25$	$\beta: 1.00 \to 1.50$	$\beta: 1.00 \to 1.75$
0.0	14.1	12.3	4.7	3.5	2.8	2.3	1.9
0.2	14.1	12.3	4.7	3.5	2.8	2.3	1.9
0.4	14.1	12.4	4.7	3.5	2.8	2.3	1.9
0.6	14.1	12.4	4.6	3.4	2.8	2.3	1.9
0.8	**14.1**	**12.4**	**4.7**	**3.4**	**2.8**	**2.3**	**1.9**
1.0	14.1	12.4	4.7	3.4	2.7	2.3	1.9

Note: This table reports cumulative lockdown intensity over the sample period, defined as $\sum_t \alpha_t$. Rows vary ω_I, while columns vary the (anticipated) path of the transmission rate β. The benchmark case keeps $\beta = 1.00$ throughout ("constant β"). In the anticipated-change scenarios, β equals 1.00 up to $T = 40$ and then changes to the value indicated by the column heading (0.25, 0.50, 0.75, 1.25, 1.50, or 1.75); after $T = 40$, β remains at the new level. All simulations start from the same initial condition and differ only in the assumed path of β.

Table 9: Cumulative Lockdown Intensity under the Two-State Markov Model: Transmission Rate (β)—Alternative Values of ω_I—

ω_I	$\beta: 1.00 \to 0.25$	$\beta: 1.00 \to 0.50$	$\beta: 1.00 \to 0.75$	constant β	$\beta: 1.00 \to 1.25$	$\beta: 1.00 \to 1.50$	$\beta: 1.00 \to 1.75$
0.0	13.8	5.4	4.2	3.5	2.6	2.0	1.6
0.2	13.8	5.4	4.2	3.5	2.6	2.0	1.5
0.4	13.9	5.4	4.2	3.5	2.5	2.0	1.5
0.6	13.9	5.5	4.1	3.4	2.5	1.9	1.5
0.8	**13.9**	**5.4**	**4.1**	**3.4**	**2.5**	**1.9**	**1.5**
1.0	14.0	5.5	4.1	3.4	2.5	1.9	1.5

Note: This table reports cumulative lockdown intensity over the sample period, defined as $\sum_t \alpha_t$. Rows vary ω_I, while columns vary assumptions about the transmission rate process for β. The benchmark case keeps $\beta = 1.00$ throughout ("constant β"). In the six two-state Markov cases, β starts at 1.00 and, from $t = 0$, can switch to an alternative absorbing state with per-period transition probability $p_2 = 0.025$. The alternative-state value is 0.25, 0.50, 0.75, 1.25, 1.50, or 1.75, as indicated by the column headings. In the realizations summarized here, the switch occurs at $T = 40$; after switching, β remains at the new level. All simulations start from the same initial condition and differ only in the assumed process for β.