

Covid-19 and Output in Japan*

Daisuke Fujii[†] Taisuke Nakata[‡]
University of Tokyo University of Tokyo

First Draft: January 2021

Abstract

Using a tractable macro-SIRD model, we explore the relationship between the spread of Covid-19 and output in Japan. We also examine the consequences of adopting alternative criteria for ending the state of emergency currently in place in Tokyo. We will update our analysis weekly at <https://Covid19OutputJapan.github.io/>.

JEL: E17, E70, I18

Keywords: Covid-19; Emergency Declaration; Forecast Evaluation; Japan; Macro-SIRD model; Vaccines.


*We thank Yasushi Iwamoto, Munehika Katayama, Sagiri Kitao, So Kubota, Keiichiro Kobayashi, Antoine Lepetit, Haruko Noguchi, Kengo Nutahara, Fumio Ohtake, and seminar participants at the University of Tokyo for comments and suggestions. We thank Takeki Sunakawa for reviewing our Matlab codes, designing our project's computational architecture, and setting up the website for weekly updates. Daisuke Fujii is supported by JSPS Grant-in-Aid for Scientific Research (KAKENHI), Project Number 17K13744. Taisuke Nakata is supported by JSPS Grant-in-Aid for Scientific Research (KAKENHI), Project Number 19K24682, the Center for Advanced Research in Finance at the University of Tokyo, and the Tokyo Marine Kagami Memorial Foundation. All errors are our own.

[†]Faculty of Economics, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo, 113-0033; Email: dfujii@e.u-tokyo.ac.jp. Research Institute of Economy, Trade and Industry (RIETI).

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

1 Introduction

Figure 1: The weekly number of new Covid-19 cases in Japan



Source: Ministry of Health, Labor, and Welfare.

As shown in Figure 1, the number of new Covid-19 cases in Japan has been on an upward trend since October 2020. The pace of the increase has accelerated in the past few weeks, having led policymakers and the public to ask how strictly they should restrain economic activities and how vigilant they would like to be against the risk of infection. Given unpredictable nature of the evolution of Covid-19, policymakers and the public will likely need to frequently reassess their policies and behaviors in coming week and months.

In this paper, we develop a tractable model of Covid-19 and the economy to assist such reassessment. We use our model to compute projections of Covid-19—conditional on various paths of output—and construct a projected relationship between Covid-19 and output in Japan. We also use our model to examine the effectiveness of alternative strategies for ending the state of emergency currently in place in Tokyo.

2 Model

Our model is formulated in a discrete-time with each period interpreted as a week. It consists of two parts: the epidemiological and economic part. The epidemiological part

is given by the following SIRD model:

$$S_{t+1} - S_t = -N_t - V_t \quad (1)$$

$$I_{t+1} - I_t = N_t - N_t^{IR} - N_t^{ID} \quad (2)$$

$$R_{t+1} - R_t = N_t^{IR} + V_t \quad (3)$$

$$D_{t+1} - D_t = N_t^{ID} \quad (4)$$

$$N_t^{IR} = \gamma_t I_t \quad (5)$$

$$N_t^{ID} = \delta_t I_t \quad (6)$$

S_t , I_t , and R_t denote the number of susceptible, infectious, and recovered persons, respectively. D_t denotes the number of cumulative deaths. N_t , N_t^{IR} , and N_t^{ID} are the number of newly infected persons, newly recovered persons, and deaths from Covid-19 between time t and time $t+1$, respectively. V_t is the number of newly vaccinated persons from time t to time $t+1$.¹ γ_t and δ_t are time-varying parameters for recovery and death rates, respectively.²

The economic part of our model is given by the following linear production function.

$$\begin{aligned} Y_t &= (1 - \alpha_t) A_t(\{\alpha_t\}_{j=0}^{t-1})(S_t + a_I I_t + R_t) \\ &:= (1 - \alpha_t) \bar{Y}_t \end{aligned} \quad (7)$$

Y_t is output and depends on (i) the number of persons adjusted for the fact that some infected persons cannot work or are less productive—given by $(S_t + a_I I_t + R_t)$ —and (ii) output per person—given by $(1 - \alpha_t) A_t(\{\alpha_t\}_{j=0}^{t-1})$. Output per person consists of two components. The first component, $(1 - \alpha_t)$, captures the reduction in output per person due to social-distancing or other measures aimed at reducing the risk of infection. The second component, A_t , is output per person that would prevail if no person takes measures against the risk of infection at time t . The dependence of A_t on the history of α_t is intended to capture possible hysteresis effects of having restrained economic activities in the past. We use \bar{Y}_t to denote the level of output that would prevail if no one restrained his or her economic activities at time t and refer to it as the *reference level of output*. Appendix A describes in detail what we intend to capture by the reference level of output as well as how we construct it.³

¹As we will discuss later, we assume that a person becomes vaccinated with a 80 percent probability after they receive two shots of vaccines.

²See Atkeson (2020) and Moll (2020) for exposition of SIR models.

³Our model abstracts from optimizing behaviors of individuals. See Eichenbaum, Rebelo, and Trabandt (2020) for an example of optimization-based macro-SIRD models, among many others.

The epidemiological part of our model is linked to the economic part through the following matching function for newly infected persons.

$$N_t = \frac{\tilde{\beta}_t}{POP_0} I_t S_t \quad (8)$$

where

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

and $POP_0 := S_0 + I_0$. $\tilde{\beta}_t$ denotes the infection rate. β_t denotes the “output-adjusted” or “raw” infection rate that would prevail in the absence of any decline in economic activity. β_t falls if people take actions that reduce the infection risk but do not directly affect their economic activities. For example, β_t falls if people wear masks or wash their hands when they return home. β_t also reflects the intrinsic infection rate of Covid-19. If a coronavirus variant with a higher infectious capacity spreads, it will appear as a higher value of β_t .

$(1 - h\alpha_t)^2$ captures the effect of a decline in economic activity on the infection rate.⁴ It is helpful to think of this term as a proxy of people’s mobility. While some mobility is necessary for households to consume and businesses to produce goods and services, it leads to interactions between susceptible and infectious persons and thus helps spread the disease. A high value of h means that the infection rate can be reduced a lot without reducing output that much. The value of h captures, among others, teleworkability of office work, abilities of restaurants to raise revenues through take-out services, or consumers’ willingness to switch from movie theaters to online streaming services.

Our model is not micro-founded, unlike many macro-epidemiological models recently developed in the economics profession. The advantages of our modelling approach are that the absence of tight cross-equation restrictions makes it easy for our model to fit the past data and solve the model quickly and allows us to conduct a broad set of policy experiments in a short amount of time. The disadvantage of our approach is that our analysis is subject to the Lucas critique. We judged that the advantages outweigh the disadvantage in our work because the main goal of our project is to provide with policy-makers and the public a back-of-the-envelope calculation of the effects of various policies in a timely manner.⁵

⁴See Alvarez, Argente, and Lippi (2020) and Farboodi, Jarosch, and Shimer (2020) for similar matching functions.

⁵Our use of a reduced-form model is similar in spirit to the widespread use of semi-structural models for policy analysis in central banks. For example, the FRB/US model—a semi-structural model of the U.S. economy—is used prominently for the analysis of forward guidance policies at the Federal Reserve, whereas DSGE models—which suffer from the so-called forward guidance puzzle—are not as frequently used as the FRB/US model. See Chung (2015) and Chung, Nakata, and Paustian (2018).

3 Data, assumptions, and identification of unobserved variables and time-varying parameters

We use data on N_t , N_t^{ID} , and Y_t to recover the paths of the model variables and time-varying parameters. N_t and N_t^{ID} are the number of new positive PCR test cases and the number of deaths due to Covid-19, respectively, from the Ministry of Health, Labour and Welfare in Japan. Y_t is based on monthly estimates of real GDP computed by the Japan Center for Economic Research.⁶

We assume the following initial conditions: $S_0 = 125.7M$, $I_0 = 1$, $R_0 = 0$, and $D_0 = 0$. Throughout the analysis, we set $\gamma_t = 7/5$ so that the effective reproduction number over the past four months is 1.12, the average effective reproduction number over the past four months reported in the Toyo Keizai website, which in turn is based on Mok Jung, Akhmetzhanov, Mizumoto, and Nishiura (2020).⁷ $V_t = 0$ for all t up to today.

3.1 Finding S_t , I_t , R_t , and D_t

Given S_0 and the paths of V_t and N_t , we can recover the path of S_t using equation (1). Given I_0 and the paths of N_t^{ID} and γ_t , we can recover the path of I_t using equation (2). Given R_0 and the path of I_t , we can recover the path of R_t using equation (3). Given D_0 and the path of N_t^{ID} , we can recover the path of D_t using equation (4).⁸

3.2 Finding δ_t , β_t , $\tilde{\beta}_t$, and α_t

Once we recover the path of D_t , we can recover the path of the death rate (δ_t) using equation (6). Once we recover the paths of I_t and S_t , we can use equation (8) to recover the path of $\tilde{\beta}_t$.

We make an assumption about the path of \bar{Y}_t , which is described in detail in Appendix A. Our assumed path grows very slowly over time and is essentially flat. Using the assumed path of \bar{Y}_t , we can recover the path of α_t using equation (7). Because of the


⁶The monthly real GDP from the Japan Center of Economic Research can be accessed here: <https://www.jcer.or.jp/en/economic-forecast/monthly-gdp>. We assume the same GDP values for all weeks in the same month. If a week spans over two different months, its weekly GDP is prorated using two values.

⁷This value of γ_t implies the expected recovery duration of 5 days. While this duration seems too short at a first glance, recall that we measure N_t by the number of new positive PCR test cases reported to the Ministry of Health, Labour and Welfare and that there are typically several days between the time when a person becomes infected and the time when that person shows symptoms and go to a doctor to take a PCR test.

⁸See Fernández-Villaverde and Jones (2020) for a similar identification strategy.

essential flatness of \bar{Y}_t , fluctuation of α_t largely inherits that of Y_t .⁹

Figure 2: Mobility and output



Source: Japan Center for Economic Research and Google.

We obtain an estimate of h by regressing the Google mobility index (M_t) on α_t .¹⁰ As shown in Figure 2, the correlation between M_t and Y_t is high, and thus the correlation between M_t and α_t is also high because \bar{Y}_t is essentially flat. Given the path of α_t and the estimated h , we can recover the path of β_t using equation (9).

The last weekly observation is the week ending in January 17th, 2021 (the seven-day period from January 11th, 2021 to January 17th, 2021).

Figures 3 and 4 show the identified paths of time-varying parameters and model variables, respectively.


4 Conditional projections of Covid-19

We use our model to compute projections of Covid-19 conditional on various paths of output. In computing these conditional projections, we make the following assumptions

⁹The most recent monthly GDP available from the Japan Center for Economic Research is for November 2020. We impute the path of output for December 2020 and the first two weeks of January using the mobility index during these periods and based on the estimated relationship between the mobility index and GDP from August 2020 to November 2020.

¹⁰Google's COVID-19 Community Mobility Reports provide mobility indices across different categories of places at a daily frequency. We pick the following four categories as we think they compose a good measure of mobility which affects the overall infection rate: parks, transit stations, retail and recreation, and workplaces. For each week, we compute the average of the median values of those four series to construct M_t . We use median values to eliminate the irregularity of holidays.

Figure 3: History of time-varying parameters



Source: Authors' calculation.

regarding the evolution of time-varying parameters. Note that, as mentioned earlier, the last weekly observation is the week ending in January 17th, 2021. Thus, the first week of projection is the week ending in January 24th, 2021.¹¹

We assume that the death rate (δ_t) and the output-adjusted infection rate (β_t) will be constant at their average values over the most recent five months.


We assume that the vaccine distribution begins in the first week of March 2021.¹² The number of vaccine shots administered will increase from zero in the last week of February 2021 to 4M in the final week of May 2021. Thereafter, 4M vaccine shots will be administered per week over the projection horizon. We put our assumption into perspective by reviewing some international experience so far in Appendix B.

We assume that a person receives the second vaccine shot 28 days after the first vaccine shot. For simplicity, we assume that a person remains susceptible between the first and second shots and that 80 percent of persons who have received two vaccine shots will obtain full immunity. With these simplifying assumptions, we obtain our baseline projection of

¹¹The results shown in this section and the section that follows were generated on January 20th, 2021.

¹²Our starting-date assumption is consistent with the statement made by Prime Minister Suga during a press conference earlier this month.

Figure 4: History of Covid-19 and output



Source: Authors' calculation, Japan Center for Economic Research, Ministry of Health, Labor, and Welfare.

Note: In the bottom-right panel, the dashed line shows the reference level of output.


V_t that increases from zero in the last week of March 2021 to 1.6M ($\approx 0.8 * 4M/2$) in the final week of June 2021. Thereafter, V_t will remain 1.6M throughout our projection horizon.

We condition our projection on various simple paths of α_t . In particular, we consider a set of paths whose initial value (the first week in the projection) is positive, declines linearly to zero in the first six months, and remains at zero thereafter. This pattern is intended to parsimoniously capture the pattern of α in the aftermath of the national emergency declaration in early April 2020.

Figure 5 shows our projection of Covid-19 conditional on three alternative paths α_t . The average output loss over the next 12 months associated with these paths are 1.5%, 2%, and 3%—shown by the red, black, and blue lines, respectively. According to the figure, suppressing the economic activity to a level in line with that observed in the aftermath of the national emergency declaration of April 2020 does a very good job in containing the spread of Covid-19, as can be seen in the blue lines. The economic decline that is half of this most severe case—shown by the red lines—would initially lead to a small transitory

decline in the number of new infections, followed by another surge which peaks around July 2021. Under the intermediate case—shown by the black lines—the number of new infections declines by about half in the first three months. As the economic activity picks up surpressingagain, the number of new infections would increase again and peaks in July 2021 at a level similar to the first two week of January 2021. Thereafter, the number of new infections declines slowly as the vaccine distribution progresses.

Figure 5: Conditional projections of Covid-19




Source: Authors' calculation.

5 Relationship between Covid-19 and the economy

5.1 Baseline Case

Figure 6: Projected relationship between Covid-19 and output



Source: Authors' calculation.

Note: The vertical axis shows the number of cumulative deaths by the end of the next 12 months. The horizontal dashed line indicates the total number of Covid-19 deaths during 2020.

The solid black line in Figure 6 shows the set of pairs of cumulative deaths by the end of the next 12 months and the average output loss over the next 12 months associated with various paths of α_t . As a reference, with our imputed Y_t for December 2020, the average output loss in 2020 is a touch below 4 percent. Recent projections for the growth rate of GDP in Japan by IMF, OECD, and World Bank are 2.3 percent, 2.3 percent, and 2.5 percent, respectively. The output loss in 2021 consistent with these growth projections is about 1.6 percent.¹³

According to the figure, our model predicts 7,000, 17,000, and 42,000 deaths by the end of the next 12 months if the average output loss over the next 12 months is 3%, 2%, and 1.5%, respectively. As discussed later, these values are substantially higher than what our model projected just two weeks ago, reflecting a sharp spike in the number of new infections a week ago (the week ending in January 10th, 2021). Although we expect the

¹³The most recent projections by OECD, IMF, and World Bank are published in October 2020, December 2020, and January 2021, respectively.

curve to shift down in coming weeks, these numbers indicate just how striking the recent spike was.

Our trade-off curve is concave, reflecting the explosive dynamics inherent in the SIRD model as well as our nonlinear specification of the matching function. One key implication of this concavity is that we can save more lives by reducing output by one unit when the output loss is expected to be small and the disease is out of control than when it is expected to be large and the disease is contained pretty well. In other words, there are *diminishing returns to scale* to reducing output.

Grey areas with varying darkness indicate the degree of uncertainty regarding the relationship between Covid-19 and output. The darkest and the second darkest grey areas indicate 20- and 40-percent confidence sets, respectively. The second lightest and the lightest grey areas indicated 60- and 80-percent confidence sets, respectively. These confidence sets are constructed as follows. We compute the standard error of the estimated h as well as the standard errors of the average values of the raw infection rates and deaths rates.¹⁴ Assuming that they are independently and normally distributed, we draw 40,000 sets of h , β , and δ . For each draw, we compute the trade-off curve. We then look at different percentiles at each level of output loss.

We highlight two features of uncertainty about the trade-off curve. First, uncertainty is higher when the expected output loss is smaller. This feature is driven by the explosive dynamics inherent in the SIRD model: a small difference in the effective reproductive number makes a larger difference in the number of cumulative deaths when the effective reproductive number is higher. Second, uncertainty is high at any level of output loss. Even at the right edge of the figure in which the output loss is 3 percent and uncertainty appears relatively low, the 80 percent confidence set ranges from 6,000 to 16,000.¹⁵

5.2 Sensitivity analysis


Figure 7 shows how deviations from our baseline specifications affects the relationship between Covid-19 and the economy.

The top-left panel of Figure 7 shows how the trade-off curve depends on the assumed path of the raw infection rate. The panel shows that a small increase in the raw infection

¹⁴Here, we are essentially saying that the raw infection rate and the death rate are randomly distributed with population means and that we estimate their population means by the sample averages over the past five months.

¹⁵Note that, in constructing the confidence set, we only consider parameter uncertainty—uncertainty that we have ways of quantifying, albeit imperfect. There is also uncertainty about the pace of vaccine distribution. Model misspecification generates another dimension of uncertainty. Thus, our confidence set should be seen as understating the true degree of uncertainty we currently face.

Figure 7: Projected relationship between Covid-19 and output:
Sensitivity Analysis



Source: Authors' calculation.

Note: The vertical axis shows the number of cumulative deaths by the end of the next 12 months.

rate increases the number of cumulative deaths by tens of thousands if the output loss is expected to be small.

The top-right panel demonstrates the importance of improving teleworkability, promoting flexible work arrangement, or encouraging households and businesses to substitute their economic activities in contact-intensive sectors with those in less contact-intensive sectors.

The bottom-left panel demonstrates how sensitive our projection is to the specification of the matching function for new infections. In the baseline projection, we assume a quadratic function for how α affects the infection rate—recall $(1 - h\alpha_t)^2$ in equation (9). The panel considers two alternative functions— $(1 - h\alpha_t)^{1.5}$ and $(1 - h\alpha_t)^{2.5}$.

The bottom-right panel demonstrates the benefit of distributing vaccines at a faster pace. One key feature of the panel is that the benefit of distributing vaccines at a pace

faster than in the baseline scenario is larger when the output loss is expected to be smaller. When the output loss is expected to be large, the disease is contained pretty well anyway so that the marginal value of a faster vaccine distribution is smaller.

These analyses demonstrate the importance of policymakers and the public to pursue policies or individually take actions to shift the trade-off curve down. Even though short-run lockdown policies are sometimes necessary to contain the spread of the disease in response to a sharp increase in the number of new infections, they do adversely affect the economy at least in the short run. Better health policies and better individual habits contain the spread of the disease without necessarily reducing output.

5.3 Evolution of the relationship

The projection of Covid-19 depends on the initial conditions as well as the projected paths of the raw infection rate and the death rate, which in turn depend on the recent realizations of them because we use the average values over the most recent five months for projection. As a result, the projected trade-off varies over time. To illustrate this point, Figure 8 compares the baseline trade-off curve shown earlier with trade-off curves computed one and two weeks earlier using the data available up to that point.

The most recent trade-off curve is a touch below the curve from one week ago and substantially above the curve from two weeks ago. Because the number of new infections was substantially lower two weeks ago than now, the five-month average values of the raw infection rate was substantially smaller back then. Accordingly, the number of cumulative deaths was substantially smaller two weeks ago than now.


5.4 Relationship between Covid-19 and output in 2020

Figure 9 shows the relationship between the number of deaths and the average output loss in 2020, based on counterfactual simulations of Covid-19—computed conditional on various counterfactual paths of α_t . In this exercise, our counterfactual simulation starts from the third week of January 2020 when the first Covid-19 case was reported.¹⁶ We find that, if the average output loss in 2020 had been 3% and 5%, instead of the actual 3.9%, the number of deaths in 2020 would have been about 38,000 and 530, respectively, instead of the realized value of 3,598.

The slope of the trade-off curve at the realized pair of deaths and the output loss—shown by the red circle—can be seen as capturing the value of statistical life, albeit with

¹⁶In this exercise, we consider alternative paths of the economy in which the path of α_t is multiplied by a constant for any time t from the last week of January 2020 to the last week of December 2020.

Figure 8: Projected relationship between Covid-19 and output:
January 17th, 2021 versus January 3rd and 10th, 2021



Source: Authors' calculation. The vertical axis shows the number of cumulative deaths by the end of the next 12 months.

Note: The vertical axis shows the number of cumulative deaths by the end of the next 12 months.

many caveats. The implied value of statistical life based on the slope of the curve in this figure is 145 years. In a future work, we plan to apply our methodology to other countries and explore cross-country heterogeneity in the implied value of statistical life based on realized health and economic outcomes in 2020.¹⁷

5.5 Discussion


There are many factors absent in our analysis but that would be of utmost importance if this type of analysis were to be used to inform the decisions of policymakers and the public. Here, we discuss two such factors—hospital capacity and suicides.¹⁸

We abstract from a potential increase in the death rate due to hospital congestion. Not only do overcrowded hospitals contribute to an increase in the death rate among Covid-19 patients, they also increase the death rate from other diseases by constraining the supply of medical resources. Taking this consideration into account would increase

¹⁷See Hall, Jones, and Klenow (2020) for discussion of the value of statistical life in the context of Covid-19.

¹⁸We plan to incorporate these factors into our analysis.

Figure 9: Relationship between Covid-19 and output in 2020



Source: Authors' calculation.

Note: The vertical axis shows the number of cumulative deaths by the end of 2020. The red filled circle indicated the realized pair of death and the average output loss in 2020.

the overall deaths associated with Covid-19—both direct and indirect—particularly when the output loss is small and the number of Covid-19 cases is large.

We also abstract from a potential increase in the number of suicides associated with prolonged economic distress. According to Chen, Choi, and Sawada (2009), the number of suicides per capita is more responsive to the unemployment rate in Japan than in other countries. The average unemployment rate in the first eleven months of 2020 is 2.8 percent, up from 2.4 percent in 2019. The number of suicides in the first eleven months of 2020 is 19,225—up from 18,675 in the first 11 months of 2019—reversing the decade-long downward trend for the first time.¹⁹ Some private-sector analysts expect the unemployment rate to edge up in 2021, which could push up the number of suicides further going forward.

¹⁹Source: <https://www.npa.go.jp/publications/statistics/safetylife/jisatsu.html>.

6 A real-time assessment of our model’s forecasting performance

If a model were to be used as a guidepost for policy, it would be important to be aware of how reliable the model is in explaining the past and predicting the future. Because we assume time-varying parameters, we fit the past data perfectly by construction. To quantify how reliable our model may be in predicting the future, we now examine a real-time forecasting performance of our model.²⁰

For each week from the first week of September 2020 to the first week of January 2021, we compute the number of new infections and new Covid-19 deaths over the next week and the next four weeks, applying the same forecasting procedure we described earlier and using the data available up to that point. In particular, at each point in time, we re-estimate h by regressing the mobility index on α available and compute the projected paths of raw-infection and death rates from their recent-five-month averages. We evaluate our model’s forecasts conditional on the realized path of α_t , as our paper focuses on conditional projections of Covid-19.


Our exercise is not fully real-time for two reasons. First, the monthly GDP data produced by the Japan Center for Economic Research is revised historically every month, but we abstract from that historical revision in this exercise. Second, we assume that the weekly GDP data—which is imputed from the monthly GDP data—becomes available after six weeks. In practice, the discrepancy between the week in which conditional projections are prepared and the week for which the most recent GDP is available depends on which week of the month the projection is prepared. For example, if one is preparing projections in the final week of December 2020, the most recent monthly GDP is from October 2020 and one needs to impute seven weeks of GDP.²¹ However, if one is preparing projections in the second week of January, the monthly GDP for November 2020 is available and one only needs to impute five weeks of GDP.

Top panels in Figure 10 show actual and forecasted outcomes for the number of new infections for one- and four-week horizons, whereas bottom panels show their differences—forecast errors. For the one-week-ahead projection, our model’s conditional forecast tracks the overall contour reasonably well but often misses actual outcomes by a large amount. For example, our model predicted about 23,000 new cases for the first week of January, 2021 in the final week of December, which is substantially smaller than about 43,000 cases

²⁰For readers unfamiliar with the idea of real-time forecast evaluation, see Faust and Wright (2009) and ? for examples of real-time forecast evaluation exercises for macroeconomic variables.

²¹See Section 3 for how we impute Y_t .

Figure 10: Real-time forecast evaluation:
New infections




Source: Authors' calculation.

we observed. For the four-week-ahead projection, our model systematically overpredicted the number of new infections for most of 2020. However, our model's four-week forecasts prepared in the first two weeks of December 2020 underpredicted the number of new infections.

Top panels in Figure 11 show actual and forecasted outcomes for the number of new Covid-19 deaths for one- and four-week horizons, whereas bottom panels show their forecast errors. For both one- and four-week horizons, our model has overpredicted the new Covid-19 deaths most of the time. For the four-week horizon, the forecast errors for October and November 2020 are sizable.

Figure 11: Real-time forecast evaluation:
New Covid-19 deaths



Source: Authors' calculation.

7 Analysis of Covid-19 and output in Tokyo

The Japanese government declared the state of emergency in Tokyo and several other regions on January 7th, 2021, in an attempt to slow the spread of Covid-19. Since then, Governor of Tokyo has asked individuals to stay at home after 8pm, restaurants to close their doors by 8pm, and offices to reduce the number of workers going to the office by 70 percent, etc. Some officials have suggested that they would like to maintain the state of emergency until the number of new infections is down to around 500 per day in Tokyo. For the week ending in January 17th, 2021, the average number of new daily infections was 1,504.

In this section, we use our model to examine the consequences of adopting alternative stopping criteria for the current state of emergency in Tokyo.²² To do so, we recalibrate our model using the numbers of new infections, new deaths, and the mobility index in Tokyo. The monthly GDP in Tokyo is constructed as follows. Indices of Tertiary Industry

²²The results shown in this section were generated on January 20th, 2021, using the data through January 17th, 2021.

Activity and Indices of Industrial Production, compiled by the Ministry of Economy, Trade and Industry (METI), report monthly output indices for many sectors in service and manufacturing industries. The Economic Census for Business Activity reports sectoral value added in each prefecture. We merge these datasets to compute the value-added weighted average of the monthly output indices, which is our monthly prefectural GDP index.

We consider three scenarios. In the first scenario, the economic activity during the state of emergency is such that the number of new daily cases reaches 500 in the eighth week. We refer to this scenario as the baseline scenario. In the second and third scenarios, the number of new daily cases reaches 500 in the fourth and twelfth week, respectively. We refer to these two scenarios as rapid- and gradual-decline scenarios. These scenarios are implemented by adjusting the level of α during the state of the emergency.


We assume that, once the emergency period ends, α declines to 4%, its average level from September 2020 to November 2020 according to the monthly GDP for Tokyo we constructed. We also assume that, if the number of new daily cases increases to 2,000 after the end of the current emergency period, there will be another emergency declaration, which would be in place until the same stopping criterion is met. This assumption is motivated by our observation that the hospital capacity constraint seems to have become a pressing concern when the number of new infections approached this threshold in the recent past. It would be useful to examine the benefit of expanding the hospital capacity and increasing the threshold for the second emergency declaration, but such exercise is outside the scope of the paper.

Finally, we assume that the pace of vaccine distribution in Tokyo is such that V_t increases from zero in the last week of March, 2021 to 160,000 in the final week of June, 2021.²³

The left panel of Figure 12 shows the paths of new infections under the baseline scenario with alternative criteria for ending the state of emergency. The solid red line shows the path when the stopping criterion is 500 per day (or 3,500 per week), whereas other thin dashed lines shows the paths under alternative criteria. One important feature of this panel is that, if the stopping criterion is sufficiently high and the current emergency period is sufficiently short-lived, there will be a second emergency declaration down the road. Otherwise, vaccines will arrive on time to avoid the second emergency declaration. The stopping criterion of 500 is slightly above the threshold value: it leads to another

²³160,000 per week in Tokyo is—in a population-adjusted term—is broadly in line with, 1.6M per week in Japan assumed in the previous few sections because population in Tokyo is roughly 10 percent of total population in Japan.

Figure 12: Alternative strategies in Tokyo:
Baseline scenario



Source: Authors' calculation.

Note: The vertical axis in the right panel shows the number of cumulative deaths by the end of the next 12 months.


emergency declaration in late May.

The right panel of Figure 12 shows pairs of cumulative deaths and the output loss associated with various stopping criteria. Note that, when the stopping criterion leads to another emergency declaration, the pair of deaths and the output loss is interior to the trade-off frontier shown in the blue color. The light green curve located in the northeast of the frontier curve consists of pairs of deaths and the output loss under stopping criteria that eventually lead to the second emergency declaration. If the stopping criterion is 450 or below, then the economy avoids the second emergency declaration and the pair of deaths and the output loss is on the trade-off frontier. The virtue of avoiding the second emergency declaration in our exercise is reminiscent of the virtue of avoiding the second wave in standard SIR models. See Moll (2020) who elucidates how a loose lockdown that avoids the second wave can save more lives than a strict lockdown that eventually leads to the second wave.

Figure 13 shows the results for the rapid-decline scenario. In this scenario, the stopping criterion of 500 again leads to another emergency declaration later on, and as a result, the death-output outcome is inside the frontier curve. Avoiding another emergency declaration by setting the stopping criterion to, say 200, would improve both health and economics outcomes.

Figure 14 shows the results for the gradual-decline scenario. In this scenario, the

Figure 13: Alternative strategies in Tokyo:
Rapid-decline scenario




Source: Authors' calculation.

Note: The vertical axis in the right panel shows the number of cumulative deaths by the end of the next 12 months.

pair of death and the output loss associated with the stopping criterion of 500 is on the trade-off frontier because it does not lead to the second emergency declaration.

Figure 14: Alternative strategies in Tokyo:
Gradual-decline scenario



Source: Authors' calculation.


Note: The vertical axis in the right panel shows the number of cumulative deaths by the end of the next 12 months.

Figure 15 show all three scenarios together. Several lessons emerge from the right

panel. First, the first-best strategy is the strategy of a rapid decline with a low stopping criterion because the rapid-decline scenario is associated with the best trade-off frontier among all scenarios and a low stopping criterion allows the economy to be on the frontier curve. However, in countries like Japan in which the government lacks authorities to impose—or is reluctant to impose—strict social-distancing measures on its citizens, a rapid decline may not be possible. The next best strategy seems to be the strategy of a gradual decline with a moderate stopping criterion, which leads to pairs of deaths and the output loss that are in the middle part of the frontier curve.

There are two types of strategies that appear inferior to the aforementioned two strategies. The first is the strategy of a rapid decline with a high stopping criterion, which puts the economy at a high risk of inducing the second emergency declaration. The second is the strategy of a gradual decline with a low stopping criterion, which is associated with pairs of deaths and the output loss being in flat regions of the trade-off curve.

Figure 15: Alternative strategies in Tokyo:
All scenarios




Source: Authors' calculation.

Note: The vertical axis in the right panel shows the number of cumulative deaths by the end of the next 12 months.

We end this section by examining the implications of alternative vaccine assumptions for the effectiveness of alternative stopping criteria. Figure 16 shows the evolution of Covid-19 and the trade-off curves—shown in the left and right panel, respectively—under the baseline scenario with the baseline and two alternative vaccine assumptions. In the first alternative vaccine assumption, the value of V_t is twice as large as that under the baseline at any t . In the second alternative vaccine assumption, the value of V_t is half of

that under the baseline at any t .

Figure 16: Alternative strategies in Tokyo:
Alternative vaccine assumptions



Source: Authors' calculation.

Note: The vertical axis in the right panel shows the number of cumulative deaths by the end of the next 12 months.

Let's first consider the first alternative vaccine assumption with the stopping criterion of 500. In this case, once the emergency status is over, the number of new cases increases and peaks in May at a level comfortably below the level that triggers the new round of the emergency period. Thus, the pair of deaths and the output loss is on the frontier curve, which is located on the southwest part of the frontier curve associated with the baseline vaccine assumption.

The frontier curve associated with the second alternative vaccine assumption is on the northeast part of the baseline frontier curve. Under this assumption and with the stopping criterion of 500, the government declares the second emergency declaration in early June. Thus, the pair of deaths and the output loss associated with the stopping criterion of 500 is interior to the frontier curve. To be on the frontier curve, the stopping criterion has to be 300 or less.

Figure 16 underscores how important it is to distribute vaccines at a faster pace. Provided that the government successfully avoids the second emergency declaration, the choice of the stopping criterion is about which point to choose on a given trade-off frontier. The optimal choice depends on various factors, including model specifications, assumptions, and one's philosophy about life and death—factors about which empirical evidence may not be able to provide definite answers and two reasonable people can disagree with

each other. In contrast, policies of distributing vaccines at a faster pace are desirable regardless of what these factors are, as better vaccine policies move the entire trade-off curve in the southwest direction in which both health and economic outcomes are better.²⁴

8 Conclusion

In this paper, we used a tractable macro-SIRD model to explore the relationship between Covid-19 and the economy in Japan. We also used our model to examine the consequences of alternative stopping criteria for ending the state of emergency in Tokyo. Because our model is simple, our analysis should be seen as that of providing a back-of-the-envelope calculation. Yet, we believe that many qualitative insights from our analysis are likely to survive in alternative more-elaborate models.

We plan to update our analysis every week. Data and codes used for our analysis are available at <https://Covid19OutputJapan.github.io/>. We will monitor the forecasting performance of our model and modify our model assumptions and specifications as we deem appropriate. We will keep track of any changes we make to our baseline assumptions and model specifications.

²⁴The same goes to policies of reducing β and increasing h . As we saw in the sensitivity analysis in Section 5.2, the economy can attain a lower number of deaths for any given output loss with lower β and higher h .

References

- ALVAREZ, F., D. ARGENTE, AND F. LIPPI (2020): “A Simple Planning Problem for Covid-19 Lockdown,” *Covid Economics, Vetted and Real-Time Papers*, 14, 1–32.
- ATKESON, A. (2020): “On Using SIR Models to Model Disease Scenarios for Covid-19,” Quarterly Review 1, Federal Reserve Bank of Minneapolis.
- CHEN, J., Y. J. CHOI, AND Y. SAWADA (2009): “How Is Suicide Different in Japan?,” *Japan and the World Economy*, 21(2), 140–150.
- CHUNG, H. (2015): “The Effects of Forward Guidance in Three Macro Models,” Feds notes, Board of Governors of the Federal Reserve System (U.S.).
- CHUNG, H., T. NAKATA, AND M. PAUSTIAN (2018): “Optimal Monetary Policy in a DSGE Model with Attenuated Forward Guidance Effects,” Feds notes, Board of Governors of the Federal Reserve System (U.S.).
- EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT (2020): “The Macroeconomics of Epidemics,” NBER Working Paper Series 26882, National Bureau of Economic Research.
- FARBOODI, M., G. JAROSCH, AND R. SHIMER (2020): “Internal and External Effects of Social Distancing in a Pandemic,” *Covid Economics, Vetted and Real-Time Papers*, 9, 25–61.
- FAUST, J., AND J. H. WRIGHT (2009): “Comparing Greenbook and Reduced Form Forecasts Using a Large Realtime Dataset,” *Journal of Business and Economic Statistics*, 27(4), 468–479.
- FERNÁNDEZ-VILLAYERDE, J., AND C. I. JONES (2020): “Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities,” NBER Working Paper Series 27128, National Bureau of Economic Research.
- HALL, R. E., C. I. JONES, AND P. J. KLENOW (2020): “Trading Off Consumption and COVID-19 Deaths,” Quarterly Review 1, Federal Reserve Bank of Minneapolis.
- MOK JUNG, S., A. R. AKHMETZHANOV, K. MIZUMOTO, AND H. NISHIURA (2020): “Real-time Estimation of the Effective Reproduction Number of COVID-19 in Japan,” *Mimeo*.
- MOLL, B. (2020): “Lockdowns in SIR Models,” Lecture note.

Appendix

A Constructing the reference level of output

Our reference level of output at time t is the level of output at time t that would prevail if people did not impose any economic restraints at time t (that is, if people conducted their economic activities as if Covid-19 suddenly and magically disappeared at time t), taking as given the fact that the economy has suffered from the Covid-19 up to time t .

The reference level of output is different from the level of output consistent with pre-crisis trend (that is, the level of output that would have prevailed if there had been no Covid-19 crisis at all) because the reference level of output reflects the possibility that the Covid-19 crisis will leave permanent effects on the level of potential output. We would like our reference level of output to reflect the permanent effect of Covid-19 crisis on output because, otherwise, our alpha would be positive even at the steady state when people are not taking any social-distancing measures.

The reference level of output is different from the potential output because it reflects the cyclical position of the economy right before the Covid-19 crisis occurred. Suppose that the output gap was 5 percent right in the first quarter of 2020 and that output declines by 5 percent in the second quarter of 2020 due to social distancing measures. We would like to think that the 5% decline in output, and an associated decline in the mobility, reduced the infection rate through α_t in this example.

We construct the reference level of output in two steps.

(i) In the first step, we construct the path of potential output. We apply the Bank of Japan's estimate of the output gap in the fourth quarter of 2019 (1.2 percent) to GDP in December 2019 to construct potential GDP in December 2019.²⁵ For January, February, and March of 2020, we apply the expected growth rate of potential GDP computed before the Covid-19 crisis by OECD (0.63 percent, annualized) to construct the potential output for these three months.²⁶

For the remaining 9 months of 2020, we let potential GDP grow at a rate of 0.21 percent (annualized) so that the average growth rate of potential GDP in 2020 is 0.35 percent, consistent with the most recent OECD estimate.²⁷ For 2021 and 2022, we let potential GDP grow at a rate of 0.24 percent and 0.21 percent, respectively, taking on board the December-2020 projection of OECD.

²⁵The BOJ's measure of the output gap can be accessed in: https://www.boj.or.jp/en/research/research_data/gap/index.htm/.

²⁶OECD: Economic Outlook No. 106 (November 2019).

²⁷OECD: Economic Outlook No. 108 (December 2020).

Let us provide some background information on why we used two sources, the Bank of Japan and OECD, in constructing the path of potential GDP. The estimates of the output gap from these two institutions for the past three years are very similar. In particular, their estimates of the output gap for 2019 are both 1.6 percent. We apply the fourth-quarter estimate of the Bank of Japan to our monthly GDP for December 2020 to compute the starting point of our monthly potential output. We rely on the BOJ's quarterly estimate because OECD computes annual estimates only. As the fourth quarter output declined nontrivially in the fourth quarter of 2019 due to the consumption tax hike at the end of September 2019, using the fourth-quarter estimate of the output gap from the BOJ is more sensible than using the average estimate of the output gap for 2019 from OECD in constructing potential GDP in December 2019. We primarily rely on OECD in constructing the path of potential GDP from 2020 to 2022 because the Bank of Japan does not provide the growth rate of potential output.

(ii) In the second step, we add the expected cyclical component of output to the potential GDP path calculated in the first step. We take the level of the output gap in January 2020 as given and assume that the output gap will stay at that level throughout our simulation horizon.

All in all, our reference level of output grows very slowly over time and is essentially flat throughout our projection horizon.


B Distribution of vaccines

In constructing the baseline path of V_t , we assume that the number of vaccine shots administered will increase from zero in the last week of February 2021 to 4M in the final week of May 2021. Thereafter, 4M vaccine shots will be administered per week over the projection horizon.

To put our assumed path of vaccine shots into perspective, let us discuss the evolution of the numbers of vaccine shots administered in G7 countries, except Japan (Canada, France, Germany, Italy, the United Kingdom, and the United States)—shown in Figure 17. Each panel in the figure shows the daily number of vaccine shots administered per day per 1,000 people over the past few weeks in each country.

If we take the population-weighted average of this statistics across the six countries on January 13, 2021—the most recent observation at the time of writing—and multiply it by 7 to translate that number into a weekly value, we get 1.5M. It appears from the variable's contour that it is likely to continue to increase in the future in all six countries considered. If it increases at the same pace as in the first few weeks for the next five

Figure 17: The daily number of vaccine shots administered per 1,000 people in G7 countries



Source: Our World in Data.

or six weeks, the pace will be about 4M per week. As discussed above, under a set of simplifying assumptions regarding the effectiveness of each one of the two vaccine shots individuals are expected to receive, 4M vaccines shots per week means the number of newly vaccinated persons is 1.6M per week ($0.8 * 4M/2$).


Our baseline path of V_t is admittedly highly judgmental. It is not clear to what extent the experiences of these countries will be informative for predicting the evolution of vaccine distribution in Japan. Our conversation with medical experts suggests that it is logistically possible to administer vaccine shots at a substantially more rapid pace than 4M per week. However, administrative hiccups can slow down the distribution process, and the cultural aversion of the Japanese people to new and unfamiliar vaccines may mean a substantially slower pace of vaccine distribution than in other countries.

We will closely monitor the experience of other countries as well as how the Japanese governments—national and local—and medical communities prepare for the upcoming vaccine distribution in Japan. We will modify our baseline assumption on vaccine distribution as we deem appropriate.

C Monthly GDP in Tokyo

Figure 18 shows the Google mobility index and our monthly estimate of GDP in Tokyo. The figure looks qualitatively similar to the one for Japan as a whole shown in Figure 2. Both mobility and output declined sharply in April 2020 and recovered gradually thereafter. The magnitude of the decline in the mobility was much larger in Tokyo than in Japan as a whole, likely reflecting the high population density of Tokyo which makes the mobility reduction imperative for the task of reducing mobility to contain the spread of Covid-19. The magnitude of the output decline in Tokyo was about the same as that in Japan, and the economic recovery in the fall of 2020 in Tokyo is not as strong as that in Japan.

Figure 18: Monthly GDP in Tokyo



Source: Authors' calculation.