Knowledge, Germs, and Output*

Shouyong Shi
Pennsylvania State University
(sus67@psu.edu)

This version: June 2020

Abstract

To capture the dynamic interactions between knowledge diffusion and disease transmission, this paper models individuals' time allocation between learning and production. Learners and those who possess the frontier knowledge (teachers) are matched randomly. The contact rate increases in the amount of time allocated to learning, and an infected individual can transmit the pathogen to a susceptible individual in contact. An infection reduces the individual's effectiveness in learning and production. A recovery from an infection does not always come with immunity. I calibrate the pathogen to Covid-19 and start the economy in a situation where knowledge diffusion will generate 1.5% annual growth in output in the next 10 years. In a short time after the unexpected arrival of the pathogen, aggregate output drops by 13% and takes 30 weeks to recover. The death toll from infections reaches 1.2% of the population. If the society can test, trace and isolate infected leaners from the learning process, the welfare gain is 1.39% of the permanent output or, equivalently, 31.6% of the first year's output. A temporary lockdown delays the peak of infections and reduces the fall in output. Moreover, the transmission of the pathogen feeds on knowledge diffusion. If the economy starts in the steady state which has less knowledge diffusion, output drops by 4.4%, the death toll is 0.66%, and the welfare gain from controlling infections is 12.2% of the first year's output.

Keywords: Knowledge; Germs; Output; Covid-19.

* Address: 502 Kern Building, Department of Economics, Pennsylvania State University, University Park, PA 16802, USA.

1. Introduction

Knowledge diffusion and disease transmission have been interwoven throughout human history. As Jared Diamond (1999) argued convincingly, the knowledge of agriculture (farmer power) spread from the origins of early civilizations together with infectious pathogens in domesticated animals. In the present time, the primary use of knowledge may be for industrial production and services instead of agriculture. This may change the sources of knowledge and infectious pathogens, but it does not change two fundamental features of their diffusion/transmission. First, knowledge and infectious pathogens are both non-rivalrous. One is a public good and the other a public bad. Second, individuals' actions can have important consequences on the transmission and the welfare of the society, because these actions may ignore the externalities in the transmission. These two features put the interactions between knowledge diffusion and disease transmission squarely in the domain of economic analysis. To a society, a central tradeoff is between economic activities and the risk of infections and deaths.

The Covid-19 pandemic has made this tradeoff more acute than ever. Despite the alarming number of infections and fatalities of Covid-19, there is a noticeable push for reopening the economy from a lockdown intended for containing the pandemic. A fast growing literature has enhanced our understanding on the economic consequences of the pandemic (see a partial review later), but it has ignored how knowledge diffusion interacts with disease transmission. In this paper, I focus on this interaction by constructing a model where individuals choose the learning intensity and hence the contact rate in the presence of an infectious pathogen. After calibrating the pathogen to Covid-19, I evaluate the importance of knowledge diffusion in the spread of the pathogen, calculate the welfare gain from controlling infections, and analyze the effects of a lockdown.

The model economy has continuous time, one unit measure of risk-neutral individuals, two knowledge levels and one pathogen. The frontier knowledge has higher productivity than the baseline knowledge. Each individual can divide a unit flow of time between production and learning the frontier knowledge. Learning can be both formal, such as going

to school, and informal such as learning on the job. A learner randomly meets a teacher who has the frontier knowledge. The meeting (contact) rate increases in the learner's effective learning intensity that is the product of the learning time and the effectiveness. An individual's disease status can be susceptible, infectious and immune. An infection reduces the effectiveness in learning and production. In a meeting, a learner successfully acquires the frontier knowledge with a positive probability. Independently of the learning outcome, an infectious individual can transmit the pathogen to a susceptible meeting partner. A recovery from an infection can come with or without immunity. An infection increases the death rate. When an individual dies, a newborn enters the economy with the baseline knowledge and draws the disease status as either susceptible or infectious. As a benchmark, I assume that there will be no vaccine for the pathogen.

I calibrate the pathogen to Covid-19. If the pathogen does not arrive, the baseline economy will grow 1.5% per year in the first ten years. In a short time after the pathogen arrives unexpectedly, aggregate output drops by 13% and takes 30 weeks to recover. The death toll from infections reaches 1.2% of the population. Individuals respond to the pathogen dynamically. Immediately after the pathogen arrives, a susceptible learner reduces the learning intensity to reduce the risk of an infection. This does not last long. As the fraction of infected people increases over time, the chance of avoiding an infection falls quickly. This reduces the expected loss of an infection relative to the gain from acquiring knowledge. And soon, susceptible learners increase the learning intensity above the level that they would choose without the pathogen. Thus, the choice of the learning intensity based on the forward-looking behavior slows down infections initially but speeds up infections later. Because the time not spent in learning is used in production, aggregate output is higher initially, but drops by more later, than in the economy without the pathogen.

To measure social welfare, I subtract from aggregate output the value of each life lost to the infection, which is calibrated in section 5.1. If the society can test, trace and isolate infected people, the society will choose to prevent infected learners from participating in the learning process. This choice reduces the death toll from infections to a miniscule level with little reduction in output, because the number of initially infected people is very

small. The welfare gain of moving from the equilibrium to the social optimum is 1.39% of permanent output or, equivalently, 31.6% of the first year's output. If it is costly to test and control infections, the welfare gain serves as a measure of the resource that should be spent on such controls. Moreover, I find that a temporary lockdown delays the peak of infections and reduces the fall in output. However, the lockdown does not have an obvious welfare gain under the assumption that a vaccine will not emerge.

The transmission of the pathogen feeds on knowledge diffusion. To illustrate this interaction, I consider an alternative starting point where the economy is in the steady state when the pathogen arrives. In this economy, the frontier knowledge has already diffused in the population. Because of the lower learning intensity, the same pathogen has significantly smaller effects than in the baseline economy. The death toll from infections is 0.66% and output drops by 4.4%. The welfare gain from controlling infections is also smaller, at 12.2% of the first year's output.

One might have asked: since knowledge diffusion is relatively slow, how can it play an important role in a fast spreading pathogen like Covid-19? The answer is that the infectious capacity of a pathogen affects the importance of learning in opposite directions. On the one hand, most infections may occur before knowledge is diffused. On the other hand, with a highly infectious pathogen, even a small change in the learning intensity can have a large effect on how fast the pathogen spreads. This is why the number of infections differs significantly when the initial state of knowledge diffusion differs.

1.1. The Literature

On knowledge diffusion, this paper follows Lucas and Moll (2014) to focus on the time allocation between learning and production as modeled by Ben-Porath (1967). Lucas and Moll (2014) analyze the balanced growth path by assuming that knowledge is unbounded and continuously distributed in the economy. In contrast, the current model has only two levels of knowledge and so growth dies down eventually. By analyzing the transition to the steady state, this paper improves the understanding of how the stage of knowledge diffusion interacts with disease transmission. Another difference from Lucas and Moll (2014) is that

contacts or meetings in the learning process are between learners and teachers. Thus, the possibility that a learner actively contacts someone with lower knowledge in Lucas and Moll (2014) does not exist in the current paper.

In epidemiology, the most commonly used model is the SIR model by Kermack and McKendrick (1927) that describes the dynamics of the population among three states: susceptible (S), infectious (I) and recovered (R). A notable extension is SEIR that adds an exposed (E) state between the susceptible state and the infectious state (see Linka et al., 2020, for a use of this model). There are also extensions allowing for the possibility that a recovery comes with temporary immunity that can be lost with a positive probability (e.g., Anderson and May, 1979). Atkeson (2020) uses the SEIR model to construct scenarios of the Covid-19 pandemic. Noticing the feature of asymptomatic transmission of Covid-19, Berger et al. (2020) advocate testing and case-dependent quarantine that control the rate of exposure. Given the contact rates, the law of motion of the distribution of individuals in the current paper is similar to that in the SEIR model. In particular, the exposed state is a contact between a learner and a teacher, with one of them being infectious and the other susceptible. Also, I allow for the possibility that a recovery comes with no immunity and hence is susceptible for repeat infections.

The SIR or SEIR model is not a good model for understanding disease transmission in the human population, because it does not have economic choices. There are two reasons why economic choices are critical. First, the contact rates are endogenous outcomes of individuals' choices, which determine the speed of transmission. Second, economic choices are necessary for assessing the economic damage of an epidemic or pandemic and for evaluating the welfare effect of policy. In particular, a model needs to take into account how individuals respond to a policy.

The largest challenge of analyzing disease transmission with economic choices is to analyze the forward-looking behavior. In this paper, the choice of the learning intensity depends on the expected gain from learning, which is captured by the value function and the associated Bellman equation. The current distribution of individuals across different

states affects individuals' choices which, in turn, affect the evolution of the distribution in the future. To calculate the expected gain from learning, a learner takes into account these changes in the future state of the economy. This forward-looking consideration affects how much and how long a learner wants to postpone the learning intensity in the time of a pandemic. The result is the non-monotonic response of the learning intensity to the arrival of the pathogen described earlier.

Another important ingredient of the current model is heterogeneity among individuals in the knowledge level. Learners are the ones who actively contact others to acquire knowledge. The heterogeneity implies that learners and teachers have different rates of being infected and transmitting the pathogen. A policy to contain and control infections is more effective when it targets learners than teachers.

The economics literature on epidemics and, especially on Covid-19, has been expanding rapidly. Greenwood et al. (2019) study an equilibrium model of HIV/AIDS infections. Keppo et al. (2020) advocate that the SIR model should incorporate individuals' choices in order to explain the spread of an epidemic, such as the swine flu and Covid-19. However, they model individuals' choices as functions of only the distribution of individuals in the population, which do not capture the forward-looking behavior discussed above. Incorporating the forward-looking behavior, Farboodi et al. (2020) demonstrate how the dependence of individuals' decisions on future values changes the dynamics of a pandemic. Alvarez et al. (2020) analyze the planning problem of a lockdown.

In a macro context, Eichenbaum et al. (2020) show that individuals cut back work and consumption in an epidemic, which deepens a recession. Acemoglu et al. (2020) emphasize heterogeneity among different groups in the risk of infections. They show that an optimal containment policy, such as a lockdown, should target old-age individuals. Also emphasizing the age heterogeneity, Brotherhood et al. (2020) construct a dynamic macro model that incorporates the choices of formal work, telework and social activities. They demonstrate the importance of testing and contact tracing infections. Kapicka and Rupert (2020) focus on heterogeneity in the employment status instead of age. They integrate an

SIR model with a search model of unemployment to study the consequence of the Covid-19 pandemic on unemployment.

This paper is obviously related to all the papers above. However, the emphasis on the interactions between knowledge diffusion and disease transmission is new. In addition, the paper shows the importance of a different type of heterogeneity – the heterogeneity in knowledge. Individuals with different knowledge levels have different infection rates, because they have different needs to learn. The socially efficient policy should prevent infected learners from participating in the learning process. The welfare evaluation of the pandemic uses a fully dynamic model with rational choices. This differs from Hall et al. (2020) who calculate the consumption equivalent of deaths from Covid-19 infections while fixing individuals' choices.

On epidemiology, this paper suggests a new measure of the infectious capacity of a pathogen. Epidemiologists use the effective reproductive number, which is the expected number of infections that a single infectious individual can cause given the current distribution of individuals across the disease status. The economics literature cited above also uses this reproductive number. However, the number is a myopic measurement because it ignores the fact that the distribution of individuals will change over time. Relevant for individuals' choices, the reproductive number should take into account how the distribution of individuals will change. I show that the "effective" reproductive number over-estimates the reproductive capacity of a pathogen and the over-estimation is substantial near the peak of infections.

2. The Model

2.1. Model Environment

Time is continuous. There is a unit measure of individuals in the economy who are risk neutral and have a rate of time preference r. An individual is endowed with L units of the flow of time that can be allocated to learning or production. L is normalized to 1 in the benchmark model but will be reduced in the analysis of a lockdown later. There are

two levels of knowledge: the baseline indexed by 0 and the frontier indexed by 1. The productivity of knowledge $a \in \{0,1\}$ is γ^a, where $\gamma > 1$. An individual with the baseline knowledge is a learner and an individual with the frontier knowledge is a teacher.

There is a pathogen of the level $b > 1$. An individual's disease status, denoted as $z \in \{b,1,0\}$, describes the pathogen level and whether the individual is immune to the pathogen b. An individual with $z = b$ is infected by the pathogen and is infectious. An individual with $z = 1$ is currently not infected with the pathogen b and is susceptible to an infection. An individual with $z = 0$ is immune to the pathogen. There is no vaccine for the pathogen b, and so only natural immunity is possible. There are two types of recovery from an infection. One is a recovery with immunity, which occurs at the rate μ and makes the individual immune permanently. The other is a recovery without immunity, which occurs at the rate ρ and changes the individual to $z = 1$. Such a recovery leaves the individual susceptible for a repeat infection.

An individual in the state (a,z) is referred to as an individual az. An individual az dies at the rate δ_z, where $\delta_b > \delta_1 = \delta_0 > 0$. The higher death rate for an infected individual than a susceptible individual will be important for the quantitative analysis. When an individual dies, a newborn enters the economy with knowledge $a = 0$ and draws a value $z \in \{1,b\}$. The probability of drawing $z = 1$ is equal to the ratio of individuals with $z = 1$ in the population to the sum of individuals with $z = 1$ and b. Thus, although knowledge cannot be inherited, pathogens can pass from one generation to the next.

The disease status affects an individual's effectiveness in learning and production. Denoted as ε_z, the effectiveness is

$$\varepsilon_z = \begin{cases} \varepsilon_0 = 1, & \text{if } z = 0 \text{ or } 1 \\ \frac{\varepsilon_0}{b}, & \text{if } z = b \end{cases}$$

The symbol ε_0 is kept for convenience. The effectiveness is assumed to be the same for learning and production. Let $s \in [0,1]$ be the amount of time that an individual spends in learning, i.e., the learning intensity. The effective learning intensity is $\ell_z = \varepsilon_z s$. An individual az with the learning intensity s produces a flow of output, $(L-s)\,\varepsilon_z \gamma^a$.

Only individuals with the baseline knowledge spend time in learning. As in Ben-

Porath (1967), learning includes both formal learning, such as going to school, and informal learning while working. With an effective learning intensity $\ell_z = \varepsilon_z s$, a learner randomly meets an individual with the frontier knowledge at the rate $\sigma(\ell_z)$, where $\sigma' > 0$, $\sigma'' < 0$, $\sigma(0) = 0$ and $\sigma'(0) = \infty$. Conditional on such a meeting, the individual acquires the frontier knowledge with the probability $\beta \in (0, 1)$. As in Lucas and Moll (2014), the meeting (contact) rate depends on an individual's own effective learning intensity but not on the teacher-learner ratio in the economy. This assumption captures the important feature that knowledge is non-rivalrous. Someone's learning does not crowd out other individuals' ability to learn. In contrast to Lucas and Moll (2014), a learner only meets a teacher and not another learner. This assumption is reasonable in the current setting where there are only two knowledge levels at any given time.

A contact can transmit pathogens regardless of which side initiates the contact and of whether the learner succeeds in learning. The pathogen is transmitted between two individuals in a meeting only if the disease status is b (infectious) for one and 1 (susceptible) for the other. In this case, the susceptible individual contracts the pathogen with a probability $\psi \in (0, 1)$.

There are two implicit assumptions in this modeling of the transmission of the pathogen: an infection is asymptomatic and a person cannot prevent others from contacting. Both assumptions can be relaxed at the cost of complicating the model. For example, one can introduce a cost to test for the infection or to reduce contacts by others.[1] I will later examine the effects of a lockdown.

The aggregate state of the economy at any time t is the distribution of individuals over (a, z), denoted as $f_t = (n_{azt})_{az=0b,1b,01,11,00,10}$, where n_{azt} is the measure of individuals with knowledge a and the disease status z at time t. Denote the measure of all individuals with knowledge a at t as

$$N_{at} \equiv \sum\nolimits_{z=0,1.b} n_{azt}. \tag{2.1}$$

[1] Toxvaerd (2020) and Krasikov and Lamba (2020) examine individuals' choice of social distancing when the pathogen transmission is asymptomatic.

Similarly, denote the measure of individuals with the disease status z as

$$M_{zt} \equiv n_{0zt} + n_{1zt}. \tag{2.2}$$

The average death rate in the population at time t is:

$$\bar{\delta}_t = \sum\nolimits_{z=1,b,0} \delta_z M_{zt}. \tag{2.3}$$

Because the population is equal to one, $\bar{\delta}_t$ is also the measure of newborns at time t. As assumed earlier, a newborn at time t draws $z \in \{1, b\}$ with the probability $\frac{M_{zt}}{M_{1t}+M_{bt}}$.

The infection rate of the pathogen depends on the composition of individuals in the economy. Denote

$$x_{1zt} \equiv \frac{n_{1zt}}{N_{1t}}, \quad x_{0zt} \equiv \frac{n_{0zt}}{N_{1t}}, \quad z \in \{b, 1, 0\} \tag{2.4}$$

In particular, x_{1bt} is the fraction of infected individuals among teachers, and x_{0bt} is the ratio of infected learners to the population of teachers. To calculate the infection rate of a teacher, I first calculate the rate at which a teacher is contacted by a learner with a disease status z. Let σ^*_{zt} denote the contact rate of a learner with the disease status z at t using the optimal learning intensity. The measure of contacts made by all such learners is $\sigma^*_{zt} n_{0zt}$. A fraction $\frac{n_{1z't}}{N_{1t}}$ of these contacts are with teachers whose disease status is z'. The rate at which a teacher with z' is contacted by a learner with z is

$$\sigma^*_{zt} n_{0zt} \frac{n_{1z't}}{N_{1t}} \times \frac{1}{n_{1z't}} = \frac{\sigma^*_{zt} n_{0zt}}{N_{1t}} = \sigma^*_{zt} x_{0zt}. \tag{2.5}$$

A teacher can be infected only if the teacher is susceptible and the learner is infectious, i.e., $z' = 1$ and $z = b$ in (2.5). Thus, a susceptible teacher is infected at the rate $\psi \sigma^*_{bt} x_{0bt}$, and the measure of such infections at t is $\psi \sigma^*_{bt} x_{0bt} n_{11t}$. Similarly, the rate at which a learner with z contacts a teacher with z' is $\sigma^*_{zt} x_{1z't}$. A susceptible learner is infected at the rate $\psi \sigma^*_{1t} x_{1bt}$, and the measure of such infections at t is $\psi \sigma^*_{1t} x_{1bt} n_{01t}$. At t, the measure of active cases of infections is M_{bt} and the measure of additional deaths caused by infections is $(\delta_b - \delta_1) M_{bt}$.

2.2. Value Functions and Optimal Decisions

Consider an individual az at time t. Denote the value function normalized by γ as v_{azt}, where the dependence on the aggregate state f_t is abbreviated as t. The value function represents the following expected utility from t onward:

$$v_{azt} = \mathbb{E}_t \left\{ \int_t^\infty \varepsilon_{z_\tau} \gamma^{a_\tau - 1} e^{-(r+\delta_z)(\tau - t)} \mathrm{d}\tau \mid (a_t, z_t) = (a, z) \right\}.$$

The term $\gamma^{a_\tau - 1}$ is the individual's productivity at τ relative to γ. The expectations are taken over the shocks and the meeting outcomes during the lifetime, while the death probability is already factored into the effective discount rate $r + \delta_z$. For a learner ($a = 0$) with the disease status z, the Bellman equation for the value function v_{0zt} is:

$$\begin{aligned} (r + \delta_z) v_{0zt} = \tfrac{\mathrm{d}}{\mathrm{d}t} v_{0zt} \quad & + I_{z=b} \left[\mu \left(v_{00t} - v_{0bt} \right) + \rho \left(v_{01t} - v_{0bt} \right) \right] \\ & + \max_{\ell \in [0, \varepsilon_z L]} \left[\left(\varepsilon_z L - \ell \right) \gamma^{-1} + \sigma \left(\ell \right) D_{zt} \right]. \end{aligned} \tag{2.6}$$

On the left-hand side is the "permanent income" of the individual. The right-hand side is the sum of capital gains and income flows. The first term is the capital gain caused by the change in the aggregate state. The second group of terms with the indicator $I_{z=b}$ arise from the two types of recoveries from an infection and they exist only if the individual is currently infected with the pathogen b, i.e., if $I_{z=b} = 1$. A recovery with immunity occurs at the rate μ, which changes the value to v_{00t}, and a recovery without immunity occurs at the rate ρ, which changes the value to v_{01t}. The last group of terms are the flow value of the time allocation, where the choice is on the effective learning intensity $\ell_z = \varepsilon_z s$. The flow of income, normalized by γ, is $(L - s) \varepsilon_z \gamma^{-1} = (\varepsilon_z L - \ell_z) \gamma^{-1}$. At the rate $\sigma(\ell_z)$, the individual contacts a teacher. The expected gain in the meeting is:

$$D_{zt} = \beta \left[v_{1zt} - v_{0zt} \right] - I_{z=1} \psi x_{1bt} \left[\beta \left(v_{11t} - v_{1bt} \right) + (1 - \beta) \left(v_{01t} - v_{0bt} \right) \right]. \tag{2.7}$$

In a meeting, the learner succeeds in learning with the probability β. The gain is $[v_{1zt} - v_{0zt}]$ if the learner does not change the disease status. Since the learner may get infected, the gain must subtract the expected loss from an infection. The learner can become infected only if he/she is susceptible, i.e., if $I_{z=1} = 1$. Conditional on a contact, a susceptible learner is infected with the probability ψx_{1bt}, where x_{1bt} is the fraction of teachers who are

infectious, as defined in (2.4). An infection reduces the gain to the learner by $[v_{11t} - v_{1bt}]$ if the learner succeeds in learning, and by $[v_{01t} - v_{0bt}]$ if the learner fails to learn.

The optimal choice ℓ satisfies the following condition of complementary slackness:[2]

$$\sigma'(\ell) D_{zt}\gamma \geq 1 \text{ and } \ell \leq \varepsilon_z L. \tag{2.8}$$

Let ℓ^*_{zt} denote the solution to this condition, and let $s^*_{zt} = \ell^*_{zt}/\varepsilon_z$. With the optimal learning intensity, the contact rate is $\sigma^*_{zt} \equiv \sigma(\ell^*_{zt})$. Note that for a susceptible learner ($z = 1$), the expected gain from learning, D_{1t}, depends on the distribution of individuals directly through x_{1bt}, and indirectly through the value function v_{11t} (see below). A susceptible learner's contact rate σ^*_{bt} inherits these two types of dependence on the distribution. For an infected learner ($z = b$), the expected gain from learning does not depend on the distribution directly, but the indirect dependence through (v_{01t}, v_{11t}) exists if a recovery from an infection has a positive probability of having no immunity ($\rho > 0$). The contact rate σ^*_{bt} inherits this indirect dependence on the distribution.

For a teacher ($a = 1$) with the disease status z, the Bellman equation for v_{1zt} is:

$$\begin{array}{rl} (r + \delta_z) v_{1zt} = & \frac{\mathrm{d}}{\mathrm{d}t} v_{1zt} + I_{z=b}\left[\mu\left(v_{10t} - v_{1bt}\right) + \rho\left(v_{11t} - v_{1bt}\right)\right] \\ & + L\varepsilon_z - I_{z=1}\psi\sigma^*_{bt} x_{0bt}\left[v_{11t} - v_{1bt}\right] \end{array} \tag{2.9}$$

In contrast to a learner, a teacher's productivity relative to γ is 1. Also, a susceptible teacher is infected at the rate $\psi\sigma^*_{bt}x_{0bt}$, where $\sigma^*_{bt}x_{0bt}$ is the rate at which the teacher is contacted by an infectious learner (see (2.5)). An infection reduces the teacher's value by $[v_{11t} - v_{1bt}]$. Thus, a susceptible teacher's value function v_{11t} depends on the distribution of individuals directly through x_{0bt}. If $\rho > 0$, an infected teacher's value function v_{1bt} depends on the distribution indirectly through v_{11t}.

An infectious learner ignores the negative externality of transmitting the pathogen to a susceptible teacher. An infectious teacher can also transmit the pathogen to a susceptible learner. However, this negative effect does not appear through a teacher's choice, because a teacher does not make a decision about whether to participate in the learning process. Similarly, although a teacher generates the benefit of passing the frontier knowledge to a

[2] I have used the assumption $\sigma'(0) = \infty$ to prove $\ell > 0$.

learner, the positive effect does not appear through a choice. Both types of externalities related to a teacher would surface if a teacher could choose the extent to which to participate in the learning process or the effort to fend off contacts by learners.

2.3. Distribution of Individuals: Epidemiology

Given the learning intensities, the distribution of individuals in the economy follows dynamics similar to those in the SEIR model in epidemiology. However, there are two noteworthy details. First, there is heterogeneity in the infection rate between learners and teachers, because only learners spend time in learning. Second, a recovery from an infection can come without immunity, which occurs at the rate ρ. Figure 1 depicts the flows of individuals between different states az other than births and deaths.

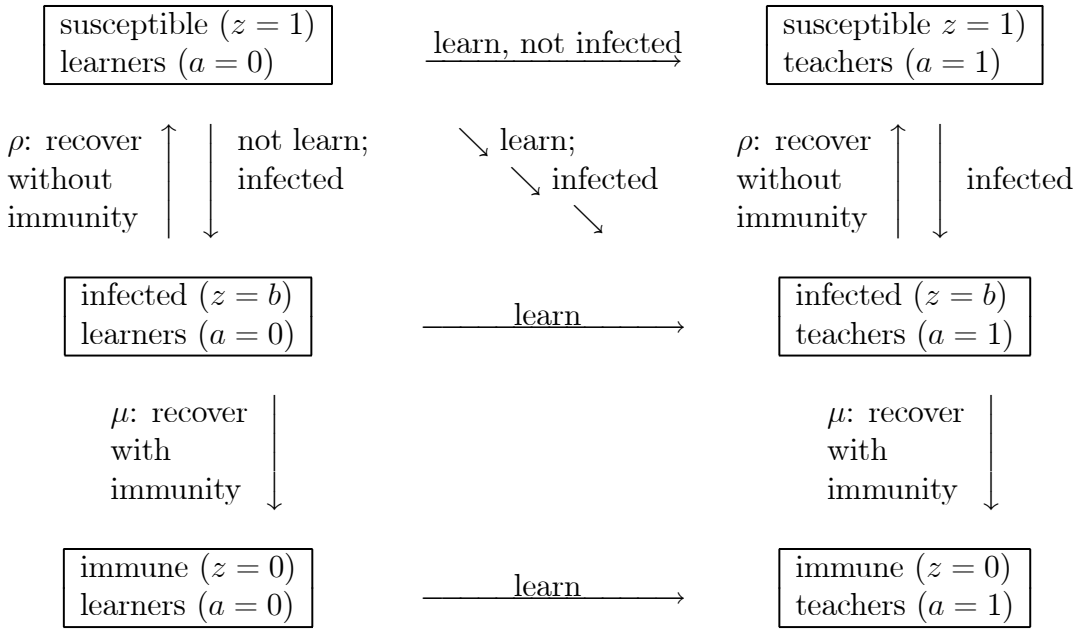

Figure 1. Flows of individuals between different states az

The group of susceptible learners ($az = 01$): The law of motion of n_{01t} is

$$\frac{\mathrm{d}}{\mathrm{d}t} n_{01t} = \frac{\bar{\delta}_t M_{1t}}{M_{bt} + M_{1t}} + \rho n_{0bt} - \left[\beta + (1-\beta)\,\psi x_{1bt}\right] \sigma_{1t}^{*} n_{01t} - \delta_1 n_{01t} \tag{2.10}$$

The first flow into the group is births. The measure of all newborns is $\bar{\delta}_t$, and a newborn has $z = 1$ with the probability $\frac{M_{1t}}{M_{bt}+M_{1t}}$. The second flow into the group is infected learners

who recover from the infection without immunity, which has the measure ρn_{0bt}. The first flow out of the group results from learning. An individual in the group meets a teacher at the rate σ^*_{1t}. If the individual succeeds in learning, which occurs with the probability β, the individual exits the group to become a teacher. If the individual fails to learn, the individual also exits the group if he/she contracts the pathogen in the meeting. The probability of becoming infected in a meeting is ψx_{1bt}. The last term in (2.10) is the outflow from the group caused by deaths.

The group of infected learners ($az = 0b$): The measure of the group, n_{0bt}, obeys

$$\frac{\mathrm{d}}{\mathrm{d}t}n_{0bt} = \frac{\bar{\delta}_t M_{bt}}{M_{bt} + M_{1t}} + \psi x_{1bt}\left(1-\beta\right)\sigma^*_{1t}n_{01t} - \beta\sigma^*_{bt}n_{0bt} - \left(\mu+\rho+\delta_b\right)n_{0bt} \tag{2.11}$$

The first term on the right-hand side is the measure of births into the group. The second term is the flow into the group from learners who fail to learn but contract the pathogen in meetings. When an individual in the group succeeds in learning, the individual exits the group. The measure of this outflow is the third term. The last term is the sum of the outflows from the group caused by the two types of recoveries and death.

Other groups of individuals: Similarly, the measures, $(n_{11t}, n_{1bt}, n_{a0t})$, obey the following laws of motion:

$$\frac{\mathrm{d}}{\mathrm{d}t}n_{11t} = \left(1-\psi x_{1bt}\right)\beta\sigma^*_{1t}n_{01t} + \rho n_{1bt} - \psi\sigma^*_{bt}x_{0bt}n_{11t} - \delta_1 n_{11t} \tag{2.12}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}n_{1bt} = \beta\sigma^*_{bt}n_{0bt} + \psi\beta\sigma^*_{1t}x_{1bt}n_{01t} + \psi\sigma^*_{bt}x_{0bt}n_{11t} - \left(\mu+\rho+\delta_b\right)n_{1bt} \tag{2.13}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}n_{a0t} = \mu n_{abt} - \delta_0 n_{a0t} + \left(1-2I_{a=0}\right)\beta\sigma^*_{0t}n_{00t}, \;\; a\in\{0,1\}. \tag{2.14}$$

The explanation for these equations is in Appendix A.

2.4. Reproductive Numbers

The effective reproductive number is the expected number of infections that an infectious individual can cause if the distribution of individuals will remain the same as the current

one. The basic reproductive number (R0) is the effective reproductive number when the pathogen just arrives. Both numbers fail to take into account the fact that the expected reproductive number should be forward-looking. Current infections change the susceptible population in the future which, in turn, will affect how many infections that a currently infectious individual can cause in the lifetime.

The reproductive numbers calculated below incorporate the effect of current infections on the distribution. For the lack of a better label, I call them the *expected reproductive numbers*. Denote the expected reproductive number as R_{abt} for an infectious individual ab and as R_t for an infectious individual on average. Examine an infectious teacher first. In a small interval of time, $\mathrm{d}t$, R_{1bt} evolves as follows:

$$R_{1bt} = \psi\sigma_{1t}^{*}x_{01t}\left(\mathrm{d}t\right) \times 1 + \left[1 - \left(\mu + \rho + \delta_b\right)\left(\mathrm{d}t\right)\right] R_{1b(t+\mathrm{d}t)}.$$

The term $\psi\sigma_{1t}^{*}x_{01t}\left(\mathrm{d}t\right)$ is the probability that a susceptible learner meets the teacher and is infected in the meeting. If the teacher recovers from the infection, with or without immunity, or if the teacher dies, the future reproductive number becomes zero. With the complementary probability, $\left[1 - \left(\mu + \rho + \delta_b\right)\left(\mathrm{d}t\right)\right]$, the teacher remains infectious, and the expected reproductive number from t+dt onward is $R_{1b(t+\mathrm{d}t)}$. In the limit $\mathrm{d}t \to 0$, the equation above becomes:

$$\left(\mu + \delta_b\right) R_{1bt} = \frac{\mathrm{d}}{\mathrm{d}t} R_{1bt} + \psi\sigma_{1t}^{*}x_{01t} - \rho R_{1bt}. \tag{2.15}$$

Similarly, for an infectious learner, the expected reproductive number R_{0bt} obeys:

$$\left(\mu + \delta_b\right) R_{0bt} = \frac{\mathrm{d}}{\mathrm{d}t} R_{0bt} + \sigma_{bt}^{*}x_{11t}\left[\psi + \beta\left(R_{1bt} - R_{0bt}\right)\right] - \rho R_{0bt}. \tag{2.16}$$

The infectious learner meets a susceptible teacher at the rate $\sigma_{bt}^{*}x_{11t}$ and infects the teacher with the probability ψ. Also, the learner succeeds in learning with the probability β, in which case his/her expected reproductive number changes to R_{1bt}. A recovery or death takes the individual out of the infectious state.

Weighting the expected reproductive number of each infectious group by the group's

relative size, the average expected reproductive number of an infectious individual is:

$$R_t = \frac{n_{0bt} R_{0bt} + n_{1bt} R_{1bt}}{n_{0bt} + n_{1bt}}. \tag{2.17}$$

The effective reproductive numbers in epidemiology can be obtained from (2.15)-(2.16) by setting $\frac{\mathrm{d}}{\mathrm{d}t} R_{1bt} = \frac{\mathrm{d}}{\mathrm{d}t} R_{0bt} = 0$. The results, denoted with a superscript e, are:

$$R^e_{1bt} = \frac{\psi \sigma^*_{1t} x_{01t}}{\mu + \delta_b + \rho}, \quad R^e_{0bt} = \frac{\sigma^*_{bt} x_{11t} \left[\psi + \beta R^e_{1bt}\right]}{\mu + \delta_b + \rho + \beta \sigma^*_{bt} x_{11t}}. \tag{2.18}$$

The average effective reproductive number, R^e_t, is:

$$R^e_t = \frac{n_{0bt} R^e_{0bt} + n_{1bt} R^e_{1bt}}{n_{0bt} + n_{1bt}}. \tag{2.19}$$

Clearly, these effective reproductive numbers vary over time as the distribution of individuals changes. Yet, the numbers are obtained by assuming that they do not change momentarily. This is a typical inconsistency in myopic measures. Similarly, the basic reproductive numbers are inconsistent because they are effective reproductive numbers at $t = 0$ when the entire population is susceptible.

The effective reproductive number does not measure the reproductivity of a pathogen effectively. Before the peak of an epidemic, the reproductive number is typically high and expected to fall over time. By ignoring this negative term, $\frac{\mathrm{d}R}{\mathrm{d}t}$, the effective reproductive number over-estimates the reproductivity of a pathogen. As shown later in section 4.2, this over-estimation is large near the peak of new infections.[3]

3. The Equilibrium

An equilibrium consists of the value functions (v_{azt}), the choices (σ^*_{zt}) and the distribution (n_{azt}), where $a \in \{0, 1\}$ and $z \in \{b, 1, 0\}$, that satisfy (i)-(ii): (i) Given the distribution, the choice σ^*_{zt} is optimal for an individual $0z$ at time t and the value functions obey (2.6) and (2.9), and (ii) The distribution satisfies (2.10)-(2.14) and $\sum_{az} n_{azt} = 1$.

[3]One may notice the similarity between (2.15) and an asset pricing equation. The left-hand side of (2.15) is akin to the permanent income of an asset. The right-hand side is the sum of capital gains and cash flows (dividends). By ignoring the capital gain, $\frac{\mathrm{d}}{\mathrm{d}t} R_{1bt}$, the effective reproductive number equates the value of the asset to the present value of the *current* cash flow. This over-estimates the value of the asset if the current cash flow is higher than the long-run level.

A critical feature of the equilibrium is the dynamic interaction between the distribution of individuals and the optimal choices. The distribution affects the value functions through the infection probabilities. In turn, the value functions affect the gains from learning and, hence, affect the optimal learning intensities. In the reverse direction, learning intensities affect the evolution of the distribution by affecting the transition of individuals between different states az.

To examine the above dynamic interaction, use (2.8) to express the optimal effective learning intensity of a learner $0z$ as $\ell_{zt}^* = \ell_z(D_{zt})$. Define

$$\Delta_{zt} = v_{1zt} - v_{0zt}.$$

This is the gain from learning when the learner does not change the disease status. Including the possibility of getting infected, the gain from learning is D_{zt} in (2.7) which I rewrite as

$$D_{zt} = \beta\Delta_{zt} - I_{z=1}\psi x_{1bt}\left[v_{11t} - v_{1bt} - (1-\beta)\left(\Delta_{1t} - \Delta_{bt}\right)\right] \tag{3.1}$$

Subtracting (2.6) from (2.9) yields:

$$\begin{array}{rl}\left(r+\delta_z\right)\Delta_{zt} = & \frac{\mathrm{d}}{\mathrm{d}t}\Delta_{zt} + I_{z=b}\left[\mu\left(\Delta_{0t} - \Delta_{bt}\right) + \rho\left(\Delta_{1t} - \Delta_{bt}\right)\right] + \varepsilon_z L\left(1-\gamma^{-1}\right) \\ & -I_{z=1}\psi\sigma_{bt}^* x_{0bt}\left[v_{11t} - v_{1bt}\right] - \max\limits_{\ell\in[0,\varepsilon_z L]}\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_{zt}\right].\end{array} \tag{3.2}$$

Examine immune individuals first, i.e., those with $z=0$. (3.1) implies $D_{0t} = \beta\Delta_{0t}$, and (3.2) implies

$$\frac{\mathrm{d}}{\mathrm{d}t}D_{0t} = \left(r+\delta_0\right)D_{0t} - \beta\varepsilon_0 L\left(1-\gamma^{-1}\right) + \beta\max_{\ell\in[0,\varepsilon_0 L]}\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_{0t}\right]. \tag{3.3}$$

The following proposition holds (see Appendix A for a proof):

Proposition 3.1. *$D_{0t} = D_0^*$ and $\ell_{0t}^* = \ell_0(D_0^*)$ for all t, where D_0^* is the unique solution to*

$$0 = \left(r+\delta_0\right)D_0^* - \beta\varepsilon_0 L\left(1-\gamma^{-1}\right) + \beta\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_0^*\right]_{\ell=\ell_0\left(D_0^*\right)} \tag{3.4}$$

Also, $v_{00t} = v_{00}^*$ and $v_{10t} = v_{10}^*$ for all t, where

$$v_{10}^* = \frac{\varepsilon_0 L}{r+\delta_0}, \quad v_{00}^* = v_{10}^* - D_0^*/\beta. \tag{3.5}$$

For an immune individual, the value function is constant over time because immunity is permanent. Moreover, for an immune teacher, the knowledge level does not change either. Since the flow value for an immune teacher is equal to the constant level of output/income, the value function is constant. As a result, the expected gain to an immune learner's success in learning is constant, and so is the optimal learning intensity. Because the income flow and the expected gain from learning are both constant, the value function of an immune learner is constant over time.

Notice that immune individuals' value functions, given by (3.5), are independent of the distribution of individuals over az. Intuitively, this result holds for an immune teacher because such an individual's value depends solely on the individual's own output/income flow. For an immune learner, the value depends on the rate at which he/she contacts a teacher. However, because knowledge is non-rivalrous, this contact rate depends only on the learner's own learning intensity and not on the teacher/learner ratio. This feature makes an immune learner's value function independent of the distribution of individuals.

For the disease status $z = b$, (3.1) implies $D_{bt} = \beta\Delta_{bt}$, and (3.2) implies

$$\begin{aligned}\tfrac{\mathrm{d}}{\mathrm{d}t}D_{bt} = \; & (r+\delta_b)\,D_{bt} - \left[\mu\left(D_0^* - D_{bt}\right) + \rho\left(\beta\Delta_{1t} - D_{bt}\right)\right] \\ & -\beta\varepsilon_b L\left(1-\gamma^{-1}\right) + \beta \max_{\ell\in[0,\varepsilon_b L]}\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_{bt}\right]\end{aligned} \tag{3.6}$$

A notable term in this equation is Δ_{1t}, which appears in (3.9) because an infected individual can recover without immunity. Such a recovery changes z to 1, after which the gain from learning without getting infected will be Δ_{1t}. As shown later, Δ_{1t} changes over time with the aggregate state, which makes D_{bt} and s_{bt} vary over t. If a recovery without immunity is not possible, i.e., if $\rho = 0$, then Δ_{1t} vanishes from (3.6). In this case, the following proposition states the features of the value functions, optimal learning intensity, and the expected gain from learning for infected individuals (see Appendix A for a proof):

Proposition 3.2. *Consider the case $\rho = 0$. Then $D_{bt} = D_b^*$ and $\ell_{bt}^* = \ell_b\left(D_b^*\right)$ for all t, where D_b^* is the unique solution to*

$$0 = \left(r+\delta_b\right)D_b^* - \mu\left(D_0^* - D_b^*\right) - \beta\varepsilon_b L\left(1-\gamma^{-1}\right) + \beta\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_b^*\right]_{\ell=\ell_b\left(D_b^*\right)}. \tag{3.7}$$

Infected individuals' value functions are $v_{0bt} = v^*_{0b}$ *and* $v_{1bt} = v^*_{1b}$ *for all* t*, where*

$$v^*_{1b} = \frac{\varepsilon_b L + \mu v^*_{10}}{r + \delta_b + \mu}, \quad v^*_{0b} = v^*_{1b} - D^*_b/\beta. \tag{3.8}$$

Moreover, $v^*_{1b} < v^*_{10}$*,* $v^*_{0b} < v^*_{00}$ *and* $D^*_b < D^*_0$*. Furthermore,* $\ell^*_0 \geq \ell^*_b$*, where the inequality is strict if both choices are interior or if both choices are at the upper corner.*

If $\rho = 0$, infected individuals only transition among themselves or into immune ones. Because these transition rates and immune individuals' value functions are independent of the distribution, all flow values and capital gains to an infected individual are independent of the distribution. Thus, an infected individual's value function is independent of the distribution. The possibility of $\rho > 0$ destroys this dependence by introducing transition of infected individuals into susceptible ones whose value functions depend on the distribution.

For the same knowledge level, an infected individual has a lower value than an immune individual. This is not surprising. An infected individual has a lower effectiveness in learning and production. Also, they have a higher death rate which implies that, if an infected learner acquires the frontier knowledge, the increase in output generated by the frontier knowledge will last for a shorter time than it does for an immune individual. Thus, the expected gain from learning is lower for an infected learner than for an immune learner, as captured by $D^*_b < D^*_0$ in Proposition 3.2. In turn, this lower gain from learning induces an infected learner to have lower effective learning intensity and a lower contact rate than an immune learner does. However, the learning intensity not adjusted by the effectiveness is not necessarily lower for an infected learner than for an immune learner.

Most of the dynamic interactions between learning and the distribution come through susceptible individuals. These interactions drive the dynamics of infections. The distribution affects susceptible individuals' value functions through two population statistics, x_{1bt} and x_{0bt}, which are defined by (2.4). x_{1bt} is the fraction of infected individuals among teachers, and x_{0bt} is the ratio of infected learners to the population of teachers. A susceptible learner is infected at the rate ψx_{1bt}, which affects the learner's expected gain from

learning as follows:

$$D_{1t} = \beta\Delta_{1t} - \psi x_{1bt}\left[v_{11t} - v_{1bt} - (1-\beta)\left(\Delta_{1t} - \Delta_{bt}\right)\right] \tag{3.9}$$

Through D_{1t}, the optimal learning intensity, $s_1\left(D_{1t}\right)$, depends on x_{1bt}. Also, x_{0bt} affects the expected gain from learning by affecting Δ_{1t}, as can be shown by setting $z = 1$ in (3.2). In contrast to Δ_{0t} and Δ_{bt}, the dynamics of Δ_{1t} are not self-contained even if $\rho = 0$. Instead, they depend on the distribution directly through (x^*_{0bt}, x^*_{1bt}) and indirectly through v_{11t}. To solve Δ_{1t} requires to solve the two value functions (v_{01t}, v_{11t}) together with the distribution n_t. Since this is not feasible analytically, I resort to the quantitative analysis.

4. Quantitative Analysis

4.1. Calibration and Computation

The contact/meeting rate has the following functional form:

$$\sigma\left(\ell\right) = \sigma_c \ell^{\eta}, \quad \eta \in (0,1)\,.$$

I calibrate the model to the weekly frequency. Table 1 lists the parameters, their values and the calibration targets. Appendix C describes the calibration procedure in detail. The value $\eta = 0.5$ comes from Lucas and Moll (2014) who, in turn, rely on the evidence on workers' wage growth over tenure. The relative productivity of the frontier knowledge to the baseline is set to $\gamma = 2$. As a reference, the average wage growth for young male workers in the U.S. in the first 8 years of career was 55% in the NLSY79 (Light, 2005). Such wage growth reflects only a part of the increase in a worker's productivity because it results only from the labor market experience. In addition, an individual's productivity in the model can increase because of formal learning, which is a part of the college premium. For the productivity loss caused by an infection, there is no direct evidence to use. A patient on the ventilator has negative productivity. However, there can be many infected ones who continue to work without knowing about the infection. I set $b = 1/0.7$ so that an infected individual loses 30% of productivity on average. As a reference, Dworsky et al. (2016) find that a permanently disabled worker experiences 30% earnings loss on average in the

second year following the injury. Since an infection of Covid-19 may not be a permanent disability, the value b is applied only when an individual is undergoing an infection.

Table 1. Parameters, values and calibration targets

Parameters	Value	Target
r	9.3827×10^{-4}	annual interest rate = 0.05
η	0.5	Lucas and Moll (2014)
γ	2	relative wage
b	$\frac{1}{0.7}$	health economics
ρ	0.0738	expected length for the two types of recovery together = 4 weeks
μ	0.1762	prob. of a recovery with immunity in the mean duration = 0.8
$\delta_0 = \delta_1$	2.9586×10^{-4}	expectancy of working = 65 years
δ_b	2.811×10^{-3}	additional death prob. from an infection = 0.01
ψ	0.5	incubation period = 2 weeks
β	0.0011	expected length to get infected if contact full time = 4 days
σ_c	3.5096	expected length to acquire knowledge if learn full time = 5 years

The pathogen is calibrated to Covid-19. First, when learning full time, the expected length of time it takes for a susceptible learner to be infected is $\frac{1}{\psi\sigma_c}$. This is set to 4 days so that the reproductive number of an infection in the first 10 weeks after the pathogen arrives is 3 on average. Second, the incubation period is two weeks and the average length of a recovery is four weeks. The length of time to be infected, the incubation period and the recovery time all fall in the range in Linka et al. (2020).[4] Because an infected individual is infectious only in the incubation period, these targets pin down $\psi = 2/4$, i.e., the probability of being infectious before a recovery. Together with the target on the expected time needed to acquire the frontier knowledge, these targets also help identify the learning probability β and the constant σ_c in the contact rate function. Third, the additional death probability from an infection is set to 0.01 to identify δ_b. The fatality rate of Covid-19 among *confirmed cases* is currently near 5.7% worldwide. The lower number

[4]A contact in this model is similar to the exposed state in the SEIR model used by Linka et al. (2020). These authors find that the expected length of time from being exposed to being infected is 2 to 6 days, and the incubation period is 3 to 18 days. They also cite a range of studies that calculate R_0 to be between 2.2 and 6.5.

used here is more realistic both because the majority of infections may not be detected and because the majority of deaths have occurred to patients who had pre-medical conditions. Finally, the probability of gaining immunity upon a recovery from an infection is set to 0.8, partly because there is no clear evidence to either support or reject the hypothesis that a recovery from Covid-19 comes with long-lasting immunity.[5] A sensitivity analysis later will consider the value 1 commonly used in the literature.

To measure economic activities, I use aggregate output (GDP) normalized by γ:

$$Y_t = \sum_{z=b,1,0} \left[n_{0zt} \left(\varepsilon_z L_t - \ell^*_{zt} \right) \gamma^{-1} + n_{1zt} \varepsilon_z L_t \right]. \tag{4.1}$$

The epidemiology part of the model is straightforward to compute. If the contact rates were exogenous, it would take only a short time to integrate the laws of motion, (2.10)-(2.14), and iterate to obtain the time-path of n. The economics part of the model can take a large time to compute. The contact rates result from individuals' choices of learning intensities which depend on the value functions that solve the forward-looking Bellman equations, (2.6) and (2.9). The value functions are functions of the aggregate state n, which contains five endogenous variables.[6] To reduce the computing time, I treat the value functions as functions of t instead of n. This enables me to solve the value functions in a feasible amount of time by forward iterations on (2.6) and (2.9). Appendix C describes the computation procedure in detail.

4.2. Baseline Results

Suppose that the pathogen has not existed up to $t = 0$. Assume that the population at $t = 0$ consists of 30% learners and 70% teachers, i.e., $n_{010} = 0.3$ and $n_{110} = 0.7$. If the pathogen never arrives, aggregate output will grow at an annual rate 1.5% on average in the next 10 years. In the steady state, the population will consist of 10.3% of learners and 89.7% of teachers.

[5] See Kirkcaldy et al. (2020) for a discussion on the hypothesis.

[6] There are six variables in n, but one of them can be eliminated by the condition that the total measure of individual is equal to one.

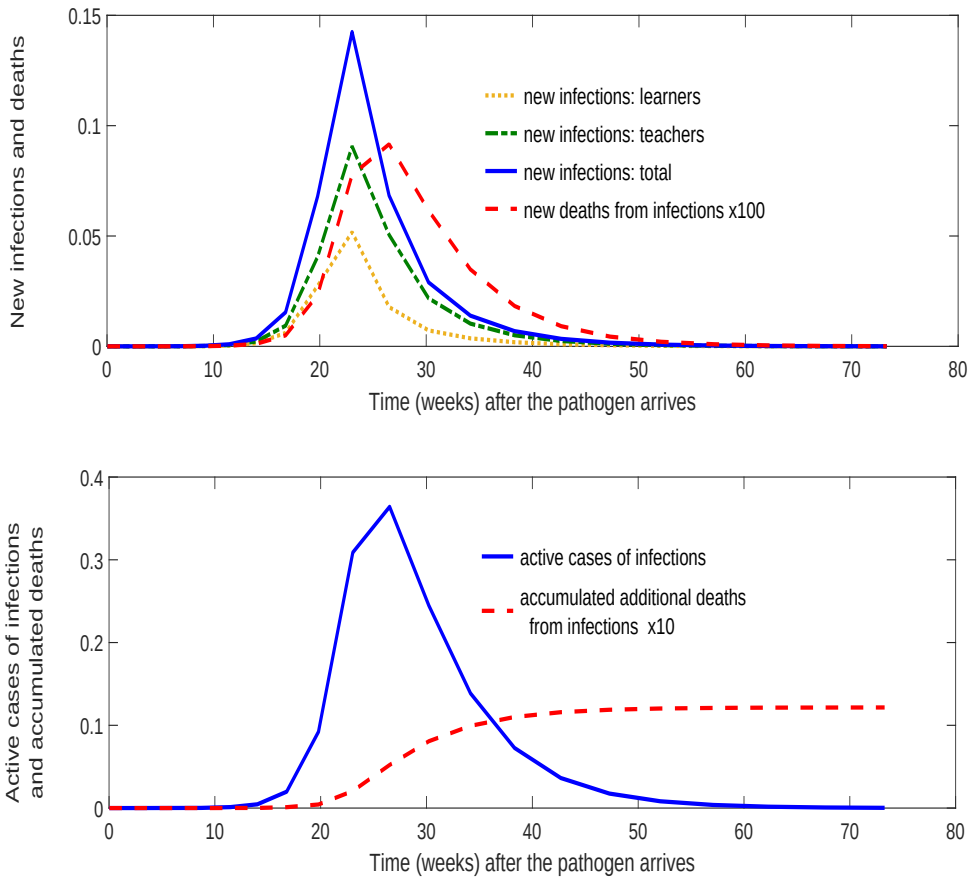

Figure 2. Infections and deaths

Now suppose that the pathogen arrives at $t = 0$ unexpectedly, with a measure of 10^{-6} learners and a measure of 10^{-6} teachers being infected. I simulate the equilibrium path of the model. Figure 2 depicts the measure of infections and additional deaths caused by infections.[7] The upper panel depicts the number of new infections among learners, $\psi\sigma^*_{1t}x_{1bt}n_{01t}$, the number of new infections among teachers, $\psi\sigma^*_{bt}x_{0bt}n_{11t}$, the sum of these two flows, and the number of new additional deaths caused by infections multiplied by 100, $100\left(\delta_b - \delta_1\right)M_{bt}$. Two features are noticeable in the upper panel in Figure 2. First, the number of new infections increases very quickly and reaches the peak in a short time (about 23 weeks). At the peak, the number of new infections is almost 14.2% of the population. The number of new infections is higher among teachers than among learners because there

[7]The model is simulated over a long time horizon, although most figures presented here report the results only for a shorter period of time.

are more teachers than learners in the economy. Second, the number of new infections is not symmetric on the two sides of the peak. It dissipates more slowly after reaching the peak than it rises to reach the peak. That is, the right tail is more protracted than the left tail. A main cause of this asymmetry is that some recoveries have no immunity and these recoveries are susceptible to repeat infections.

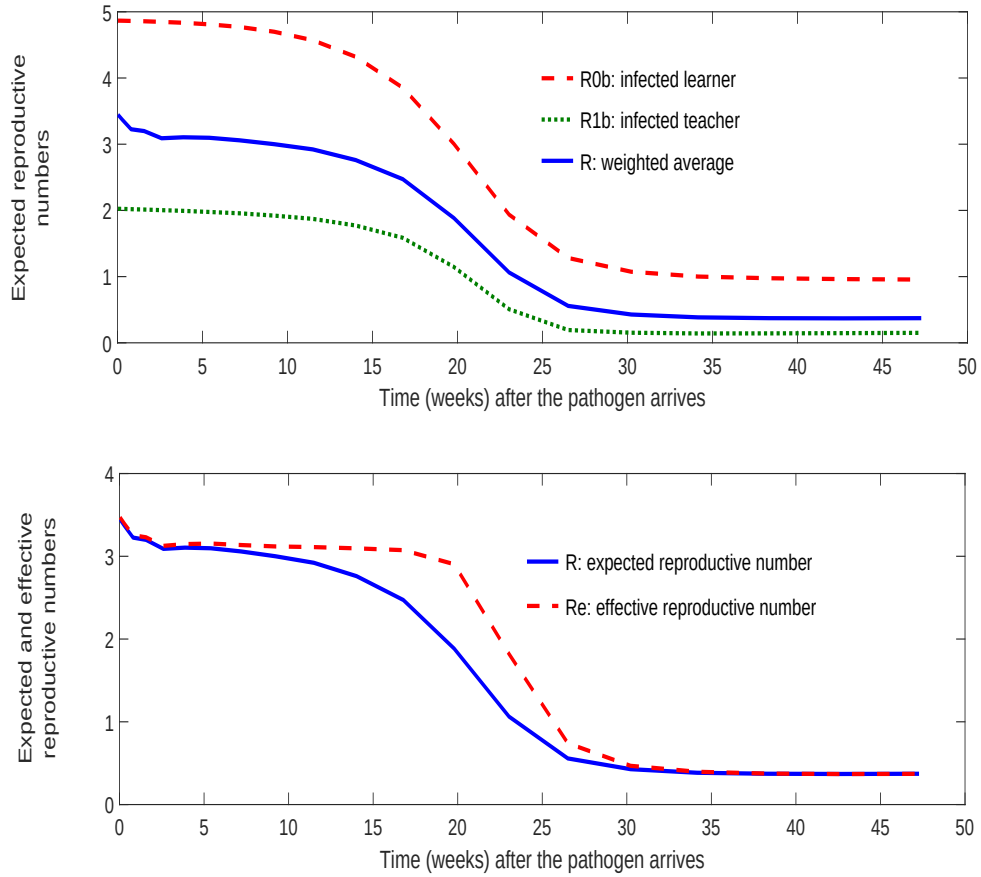

Figure 3. Reproductive numbers of infections

The lower panel in Figure 2 depicts the currently active cases of infections and the accumulated additional deaths caused by infections multiplied by 10. Currently active cases of infections include all infections up to t that have not recovered but exclude those infections that have recovered before t. It is measured by M_{bt}. This measure reaches the peak 26.5 weeks after the pathogen arrives, when 36.4% of the population are having infections. Active cases peak later than new infections because they are a stock instead of a flow of infections. It takes time for the flow to build up the stock. The accumulated additional deaths are all additional deaths caused by infections up to t, i.e., the integral

of $(\delta_b - \delta_1) M_{b\tau}$ for τ from 0 to t. This accumulated measure increases steadily around the peak of active cases of infections. After active cases of infections start to fall, the accumulated additional deaths approach the maximum asymptotically, which is 1.21% of the population.

Figure 3 depicts the reproductive numbers of the pathogen. The upper panel shows the expected reproductive numbers of an infected learner and an infected teacher, and the weighted average of the two. Before week 10, the expected reproductive numbers are relatively flat, and the weighted average in the first 10 weeks is near 3. The expected reproductive number is more than twice as high for an infected learner as for an infected teacher, because the rate of successfully learning in each contact is low. As an infected learner continues to make contacts with teachers, he/she infects the teachers at a high rate. In week 14, the expected reproductive numbers begin to fall sharply and, by week 26, the average of the expected reproductive number is below 1. Notice that the beginning of the sharp decline occurs about 8 weeks before the measure of new infections peaks. This is because the expected reproductive number is a forward-looking measure: it incorporates the expected decline before the decline happens.

The lower panel in Figure 3 contrasts the average expected reproductive number to the average effective reproductive number. Before week 10, the two numbers are very close to each other. However, the effective reproductive number is significantly higher than the expected reproductive number around the peak of new infections, i.e., between weeks 14 and 25. For example, in week 20, the expected reproductive number is 1.9 but the effective reproductive number is 2.9. Also, the effective reproductive number starts to fall significantly only after week 20, six weeks after the the expected reproductive number starts to fall significantly. These differences arise from the fact that the effective reproductive number is a myopic measure that fails to take into account the expected change in the number itself, as discussed in section 2.4.

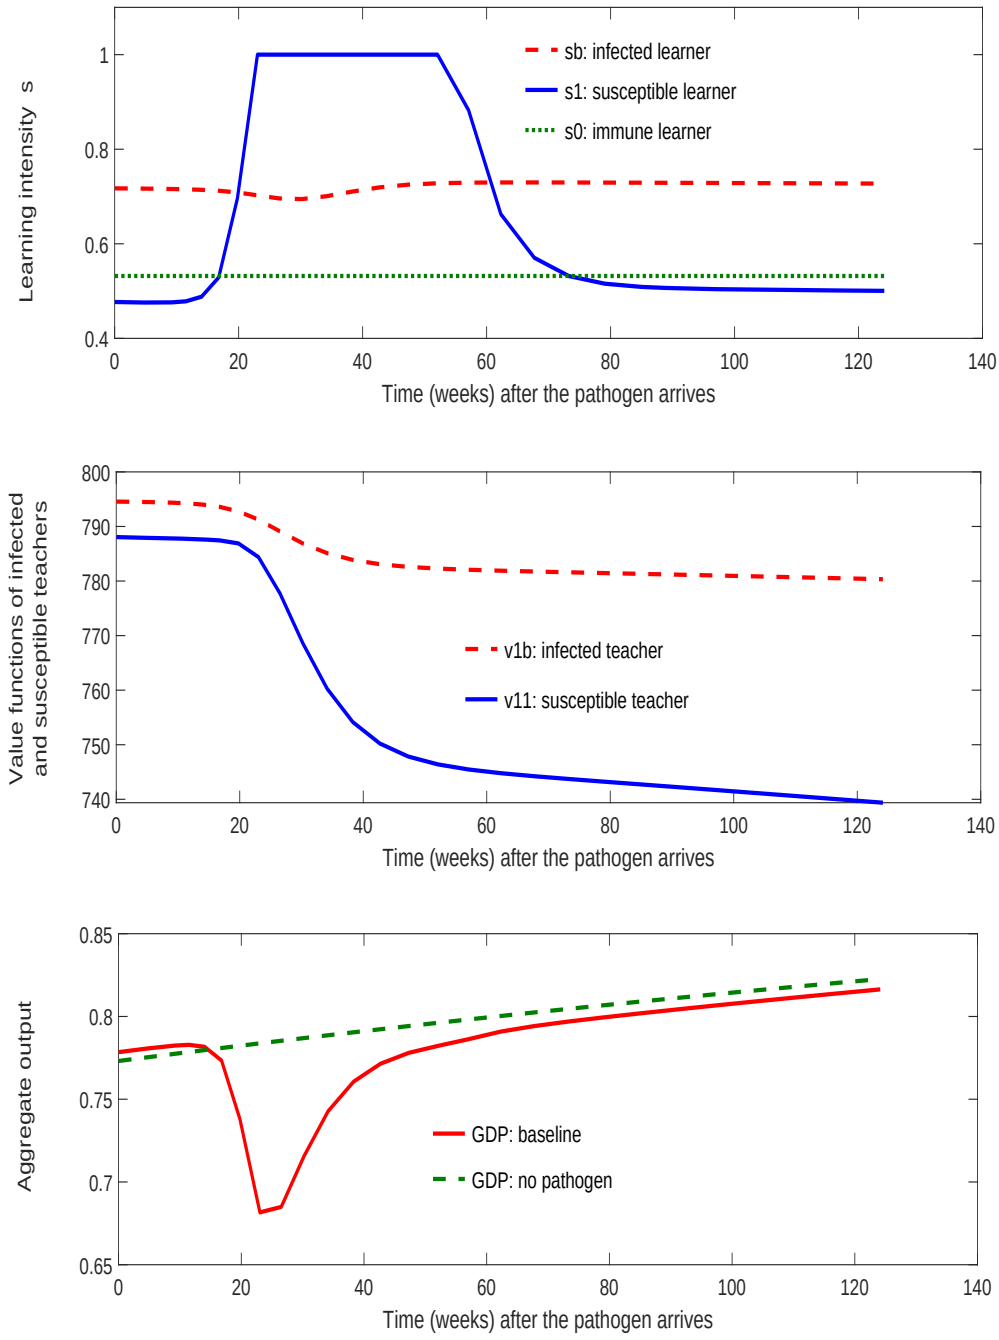

Figure 4. Learning intensities, value functions and aggregate output

To understand the importance of decisions in the dynamics of infections, I depict optimal learning intensities in the upper panel in Figure 4. An immune learner's learning intensity is constant over time. An infected learner's intensity varies over time because

$\rho > 0$, but the variation is not large. In contrast, a susceptible learner's intensity has large swings over time. After the pathogen arrives, a susceptible learner reduces intensity below the long-run level in order to reduce the chance of being infected. This low intensity does not last long. Starting in week 12, a susceptible learner's intensity increases quickly and reaches the upper bound 1 in week 23. This learning intensity stays at the upper bound for a considerable amount of time even after new infections have decreased. After week 52, a susceptible learner's learning intensity starts to fall quickly toward the long-run level.

Why does a susceptible learner's intensity have these large swings? To find the answer, let us examine the learner's gain from learning, given by (2.7) for $z = 1$. This gain has three factors. One is the gain from acquiring the frontier knowledge without being infected, $(v_{11t} - v_{01t})$. The second is the rate of being infected, ψx_{1bt}. The third is the loss from an infection, which is $\beta (v_{11t} - v_{1bt})$ if the learner succeeds in learning and $(1 - \beta)(v_{01t} - v_{0bt})$ if the learner fails to learn. The change in this loss is the dominant force driving the behavior of a susceptible learner's intensity between week 12 and week 52. During this time, the values of a susceptible learner and a susceptible teacher both fall relative to an infected counterpart's value. As depicted in the middle panel in Figure 4, this reduction is particularly significant in $[v_{11t} - v_{1bt}]$. The loss from an infection falls because the chance of avoiding an infection falls greatly.[8] As a result, the gain from learning increases, and so does a susceptible learner's intensity. The high learning intensity increases the number of contacts, which cause the increase in infections that individuals have rationally expected.

Learning intensities directly affect the dynamics of aggregate output depicted in the bottom panel of Figure 4. For a reference, the panel also depicts aggregate output when the pathogen never arrives. Immediately after the arrival of the pathogen, aggregate output is higher than the one without the pathogen, because learners scale back learning which results in an increase in production. Aggregate output peaks temporarily in week 12. After that, output drops precipitously and reaches the trough in week 23, the time when the measure of new infections peaks (see Figure 2). The decline in output in the 11 weeks

[8]The model assumes that the only way to become immune is to be infected first. This assumption and the high value of μ generate $v_{1bt} > v_{11t}$.

is 13%. Although output increases quickly after week 23, it does not reach the previous peak until week 53. The steady state of the economy is the same as the one without the pathogen, because the pathogen eventually dies out as a sufficiently large fraction of population acquires immunity.

4.3. The Effects of Reinfections

In the baseline calibration, some recovered individuals can be re-infected. To see the effects of reinfections, suppose that all recoveries come with immunity, i.e., $\rho = 0$. I recalibrate the model by keeping all other targets. Specifically, to keep the mean duration of an infection at the target of 4 weeks, the recalibration under $\rho = 0$ requires the rate of a recovery with immunity, μ, to increase from the baseline value 0.1762 to 0.25.

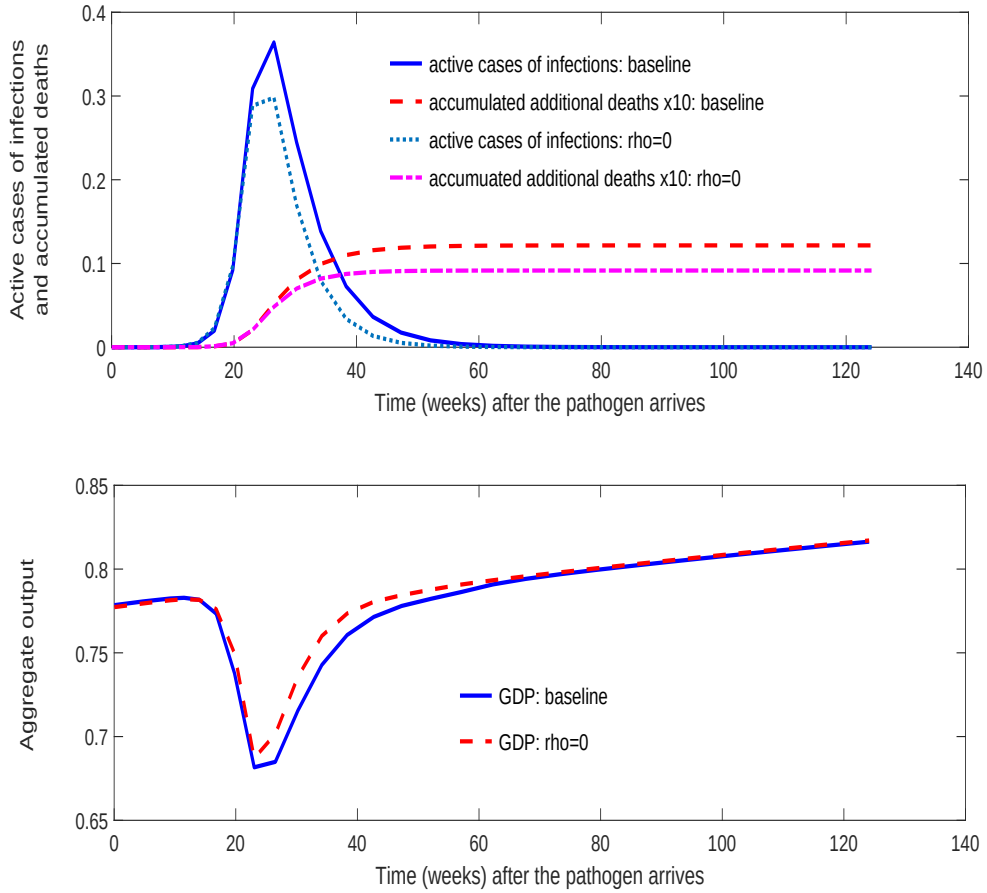

Figure 5. Recoveries without immunity: $\rho > 0$ (baseline) vs. $\rho = 0$

Figure 5 compares the dynamics of this economy with $\rho = 0$ and the baseline economy.

The upper panel shows that active cases of infections peak around the same time as in the baseline economy. The peak is lower, at 30% instead of 36.4%. After the peak, active infections are systematically lower with $\rho = 0$ than in the baseline economy. Another way to see this difference is through the accumulated additional deaths from infections. Before active infections peak, the accumulated deaths are similar in the two economies. After the peak, the gap between the two widens. Eventually, the accumulated additional deaths are 0.916% with $\rho = 0$, in contrast to 1.21% in the baseline economy. This difference, 0.294%, is big. As a reference, deaths from Covid-19 in the world are currently 0.0055% of the population.

The lower panel in Figure 5 compares aggregate output between the two economies. Because learning intensities are initially higher in the economy with $\rho = 0$ than in the baseline economy, less time is spent in production, and so aggregate output is initially lower slightly. With $\rho = 0$, the reduction in output from the peak to the trough is smaller. As a result, it takes a shorter time for output to recover with $\rho = 0$ than in the baseline economy.

4.4. Pathogen Transmission Feeds on Knowledge Diffusion

To illustrate the interaction between knowledge diffusion and the transmission of the pathogen, consider a different starting point of the economy at $t = 0$ – the steady state. In this economy, learning still occurs because newborns start with the baseline knowledge. Relative to the baseline economy, learning is much lower, since a larger fraction of the population has already acquired the frontier knowledge. Figure 6 depicts the results. As the upper panel shows, active cases of infections peak in week 79, much later than in the baseline economy. The peak is also much lower, at 10% instead of 36.4%. Low infections result in lower additional deaths from infections. These deaths approach 0.66% in the long run, in contrast to 1.21% in the baseline economy. Also, aggregate output falls by less, as shown in the lower panel in Figure 6. From the peak to the trough, output drops by 4.4% in contrast to 13% in the baseline economy. The pathogen is the same as before, but people's choices are different. The lower learning activity in the alternative economy

makes the same pathogen less infectious and the damage less severe.[9]

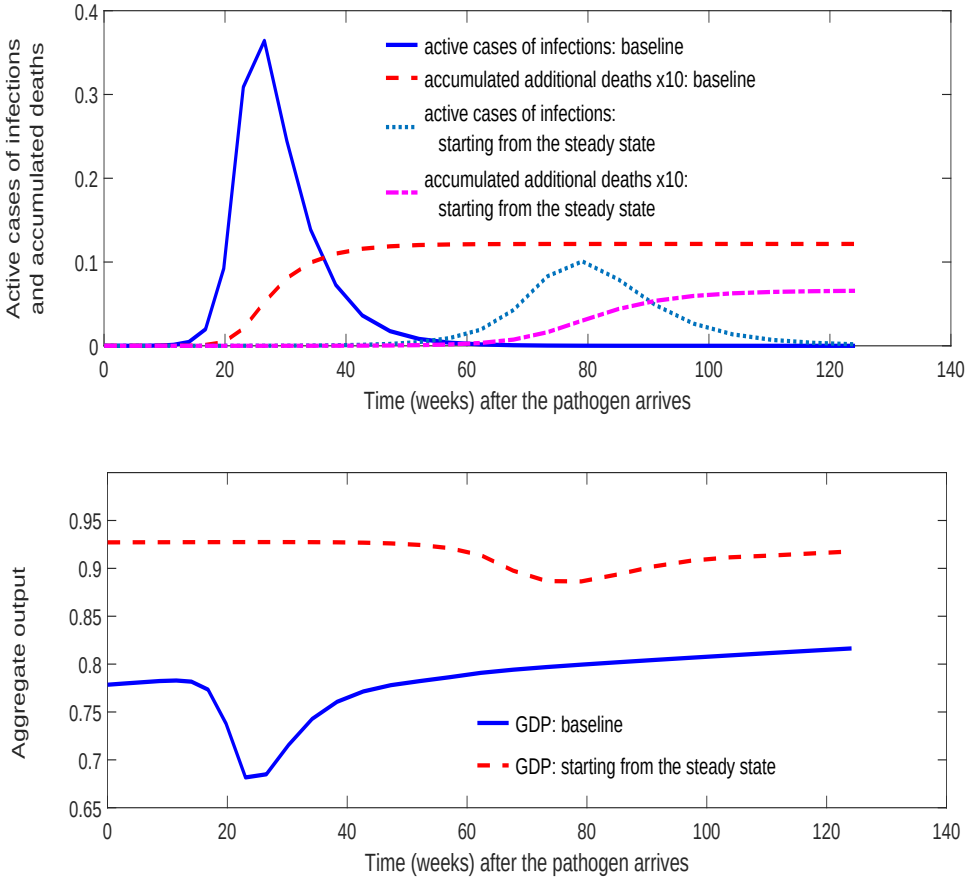

Figure 6. Effects when the economy starts in the steady state

5. Social Welfare and A Lockdown

This section accomplishes two tasks. First, it characterizes the socially efficient allocation under the assumption that the planner can identify individuals by their states az. Using the characterization as a reference, I answer the following question: if a society cannot identify individuals' disease status without a cost, how much resource should the society allocate to control infections, such as testing and contact tracing? Second, I analyze the effects of a temporary lockdown in the equilibrium.

[9]The expected reproductive number before the peak of infections is 1.54 on average.

5.1. Social Optimum

One measure of social welfare is the discounted sum of aggregate output, where output at t normalized by γ is Y_t in (4.1). This measure ignores the lives lost to infections, $(\delta_b - \delta_1) M_{bt}$. Output adjusted for these losses of lives is:

$$\tilde{Y}_t = Y_t - \zeta (\delta_b - \delta_1) M_{bt} \tag{5.1}$$

where $\zeta \geq 0$ is calibrated later. At time t, social welfare is:[10]

$$U_t \equiv \int_t^\infty \tilde{Y}_\tau e^{-r\tau} \mathrm{d}\tau, \quad i = 1, 2 \tag{5.2}$$

The permanent flow of welfare associated with U_t is rU_t.[11] Notice that this welfare measure varies with t because the allocation is not stationary.

Consider first the case where the social planner can identify individuals by their states az. The planner chooses the allocation $\{\ell_{zt}\}_{t\geq 0}$ for each z and the implied $\{n_{azt}\}_{t>0}$ for each az. The constraints are the laws of motion, (2.10)-(2.14), and the resource constraint for each z, $0 \leq \ell_{zt} \leq \varepsilon_z L$. Denote w_{azt} as the social marginal value of n_{azt}. This is the multiplier of the law of motion of n_{azt} in the planner's problem. Appendix B sets up the Hamiltonian of this problem and characterizes the socially efficient allocation. The efficient learning intensity satisfies a similar condition to (2.8):

$$\sigma' (\ell_{zt}) \gamma Q_{zt} \geq 1 \text{ and } \ell_{zt} \leq L \tag{5.3}$$

where Q_{zt} is the social gain from learning per individual $0z$, measured in output at t normalized by γ. Q_{1t} and Q_{0t} have similar expressions to D_{1t} and D_{0t} in (2.7), except that the social marginal value w_{azt} replaces the private value v_{azt}. However, Q_{bt} differs from D_{1t}:

$$Q_{bt} = (w_{1bt} - w_{0bt}) \beta - \psi x_{11t} [w_{11t} - w_{1bt}] .$$

[10] Although the discount rate in (5.2) is the rate of time preference, the death rate does matter for effective discounting, as shown by the equations for the social marginal value of n (see Appendix B).

[11] Because individuals are risk neutral in the model, this welfare measure is the same tpe that Lucas (1988) used for assessing the welfare cost of business cycles.

The last term, ,which does not appear in D_{1t}, is the negative externality that an infected learner creates on a susceptible teacher through infection. Moreover, the optimality condition of n_{azt} differs from the equilibrium counterpart by taking into account the effects of n_{azt} on births and deaths (see Appendix B for the detail).

I compute the socially efficient allocation for two values of ζ. One is $\zeta = 0$. The other is $\zeta = 970$, which is equivalent to 49.3 years of full-time output in this model. This value comes from the estimate that the average life expectancy of those who die from a Covid-19 infection is 14.5 years and the estimate by the U.S. Environmental Protection Agency (2018) that the value of a lost life is \$7.4 million in 2006 in the group age 25 to 55.[12] Since the average life expectancy in the group age 25 to 55 is 40 years but the average life expectancy of those who die from Covid-19 is 14.5 years, the EPA's estimate implies that the value of a life lost to Covid-19 is $\$7400000 \times 14.5/40 = \$2682500$. Dividing by GDP per capita in the US in 2006, \$46300, the value is 58 years of output. In the baseline model, because the average output flow per worker in the first 15 years after the pathogen arrives is 0.85, instead of 1, the value of a life lost to Covid-19 is $58 \times 0.85 = 49.3$ years of full-time output (with the frontier productivity). Then,

$$\zeta = \int_0^{49.3\times 52} e^{-rt} \mathrm{d}t = 970.$$

With both $\zeta = 0$ and $\zeta = 970$, the efficient allocation is to shut down infected learners' learning, i.e., to set $\ell_{bt} = 0$ for all t. Infected teachers are still present in the learning process and contacted by some learners, because the model does not allow the planner to shield learners from infected teachers.[13] Despite these contacts, infections dissipate quickly toward zero rather than take off as in the equilibrium. One reason for this result is that an infected teacher has a significantly lower reproductive number, as discussed for Figure 3. More importantly, as soon as a learner is infected during learning, the planner shuts the learner out of the learning process to prevent him/her from infecting others. As a result, the efficient allocation leads to only a small number of deaths, in the magnitude of 10^{-6}.

[12] Hall et al. (2020) use the same two estimates but perform a different calculation.

[13] It would also be socially efficient for the planner to prevent infected teachers from being contacted by learners if this were a choice.

Therefore, aggregate output in the efficient allocation with the pathogen is very close to that depicted in the bottom panel in Figure 4 for the equilibrium without the pathogen.

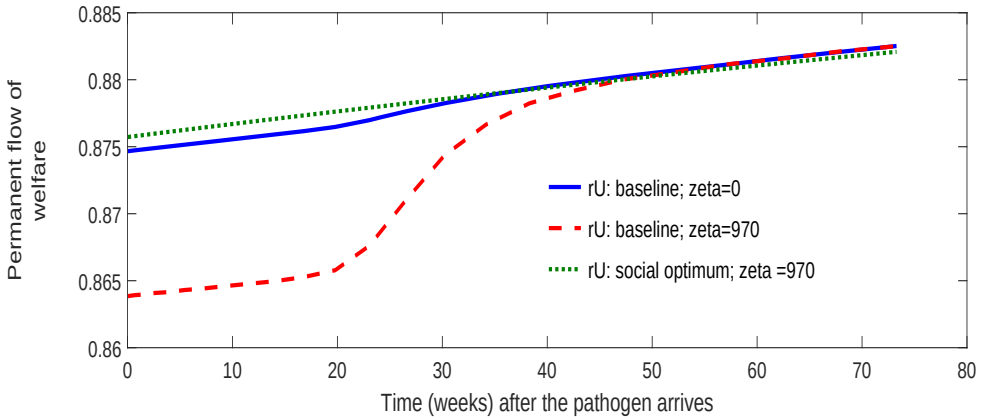

Figure 7. Welfare gains from the efficient allocation

Now consider the case where the planner cannot identify an individual by the disease status z without a cost. How much resource should the society spend to identify and trace infected individuals? To answer this question, I calculate the welfare gain from the efficient allocation above relative to the equilibrium in section 4.2. Figure 7 depicts the permanent flow of welfare rU in the equilibrium, for $\zeta = 0$ and $\zeta = 970$, and in the social optimum, for $\zeta = 970$. Social welfare in the social optimum for $\zeta = 0$ is very close to that for $\zeta = 970$, because the measure of deaths from infections is very small. With $\zeta = 0$, the equilibrium and the social optimum have similar welfare. This is the case despite that aggregate output follows significantly different dynamics in the equilibrium from the social optimum around the peak of infections (see Figure 4). This difference lasts for a relatively short period of time (less than 50 weeks), and so it does not have a significant effect on the permanent flow of aggregate output.

If $\zeta = 970$, the welfare cost of infections is substantially higher. The permanent welfare flow, rU_0, is equal to 0.8638 in the equilibrium and to 0.8757 in the social optimum. The difference is 1.36%. That is, looking forward from $t = 0$, the society is willing to spend 1.39% of the permanent welfare flow to control infections in order to achieve the efficient allocation. Since the output flow in the first year is 0.7785, this welfare gain is the following

fraction of the first year's output: $\frac{0.8757-0.8638}{0.7785\times 52\times r} = 31.6\%$. Over time as more individuals become infected in the equilibrium, the welfare gain from controlling infections falls. In week 50 and afterward, this welfare gain becomes indiscernible.

This welfare gain above is comparable with what Hall et al. (2020) obtained. These authors estimate that the society is willing to give up 41% of the first-year consumption to avoid the deaths caused by Covid-19. With 3/2 as the ratio of output to consumption in the U.S., the welfare gain calculated above in the current model is 47% of the first-year consumption. Despite the similarity in the estimate, the two models are quite different from each other. In particular, Hall et al. (2020) do not model individuals' choices: they simply calculate the consumption equivalent of the risk of death from Covid-19 infections.

The welfare gain from controlling infections is smaller if there is less knowledge diffusion. To illustrate, suppose that the economy starts in the steady state as in section 4.4. With $\zeta = 970$, the permanent welfare flow, rU, is equal to 0.9195 in the equilibrium and to 0.925 in the social optimum. By reducing infections and deaths, the social optimum increases the permanent welfare flow by 0.6%. Since the output flow in the first year is 0.925, the welfare gain is equivalent to 12.2% of aggregate output in the first year. This gain is less than 40% of the gain in the baseline economy. Thus, the higher is the growth that an economy expects to have from knowledge diffusion, the larger is the loss that the economy suffers from infections and deaths from the pathogen.

5.2. Effects of a Lockdown

A lockdown reduces the contact efficiency and reduces the amount of time an individual can spend in the market. To capture these effects, I model a lockdown as follows:

$$\text{for } t_1 \leq t < t_2\text{: } \sigma_{ct} = \sigma_{c0}\left(1 - d_\sigma\right),\ L_t = 1 - d_L. \tag{5.4}$$

The lockdown starts at t_1 and the economy reopens at t_2. The learning efficiency falls from the baseline level σ_{c0} by a proportion d_σ, and the amount of time that can be spent in learning and production falls by a proportion d_L. I set

$$t_1 = 17,\ t_2 = 16,\ \ d_\sigma = 0.5,\ d_L = 0.1.$$

The lockdown period is the one in which most infections occur in the baseline economy without a lockdown. The length of the lockdown is longer than those implemented in some countries for Covid-19. The value of σ_c is only suggestive. The value of d_L approximates the increase in the unemployment rate in the U.S. during the lockdown.[14] I assume that the lockdown is anticipated at time 0.[15]

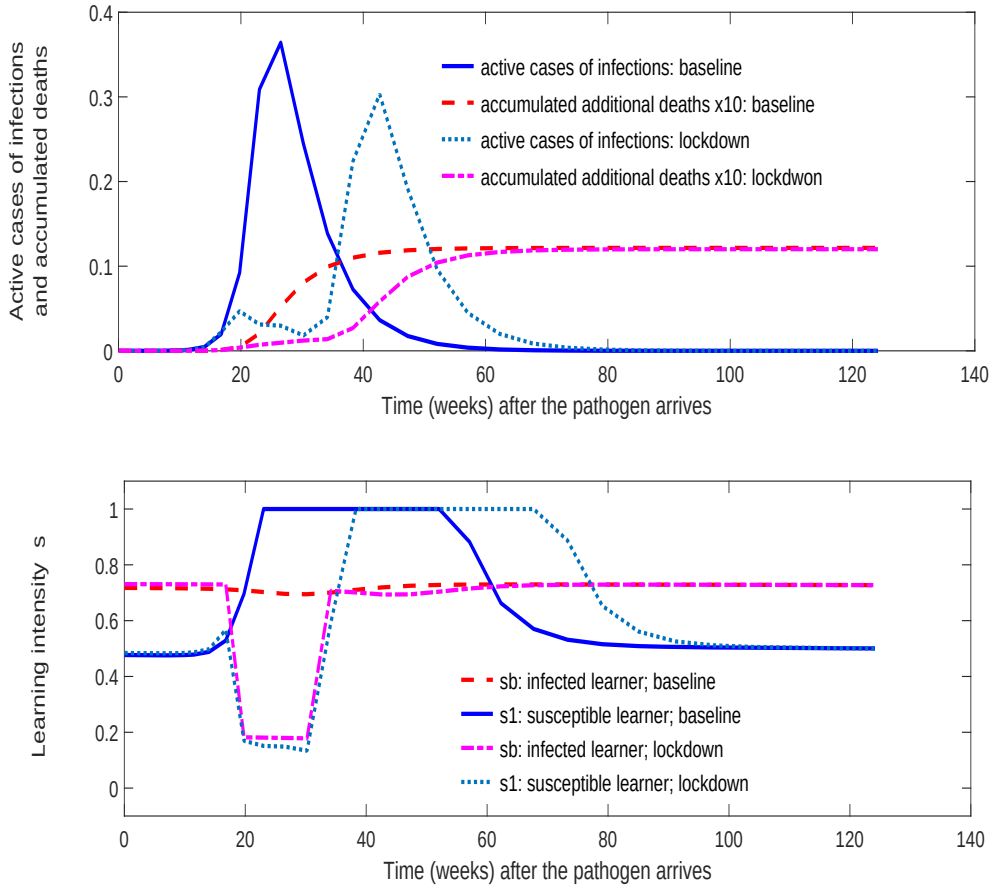

Figure 8. Effects of a lockdown on infections, deaths and learning intensity

Figure 8 depicts the effects of the lockdown. The upper panel shows that the lockdown suppresses the active cases of infections and the accumulated additional deaths to substantially lower levels while the lockdown is in effect. When the economy reopens 16 weeks later, the measure of active cases of infections peaks in week 43. The peak of active cases

[14]The reductions in both σ_c and L are necessary. If only σ_c falls but L is constant, a lockdown leads to a temporary boom as individuals shift time to production.

[15]With the lockdown, the value functions of immune individuals are no longer constant over time, but they continue to be independent of the distribution.

of infections is 30%, which is lower than the peak in the baseline economy (36.4%). The cause of the lower peak is that some individuals have been infected during the lockdown and have recovered with immunity. They do not infect others once the economy reopens. The accumulated additional deaths are also significantly lower during the lockdown than without a lockdown, but they eventually approach the similar level.

The lower panel in Figure 8 depicts learning intensities of an infectious and a susceptible learner. Before week 17, these intensities are close to the levels in the economy without the lockdown. The lockdown reduces learning intensities by more than 70%. The suppressed learning intensities keep infections low, as depicted in the upper panel. After the economy reopens, learning intensities quickly approach the high levels without the lockdown. They also stay at those high levels for roughly the same length of time as without the lockdown, before declining sharply toward the long-run level.

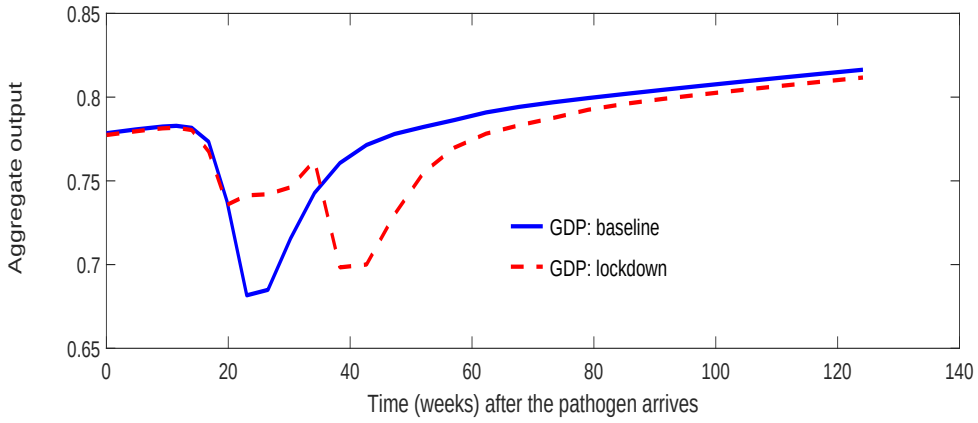

Figure 9. The effects of a lockdown on output

Figure 9 depicts the effects of a lockdown on aggregate output. The upper panel depicts the dynamics of output with and without a lockdown. Anticipating the lockdown, output is slightly lower than without a lockdown, because learners have slightly higher learning intensities that take time away from production. During the lockdown, output is on average 4.6% lower than the previous peak, smaller than the 13% drop without the lockdown.[16] After the economy reopens, output dips quickly after learning intensities and infections

[16] After the initial fall, output increases as the lockdown continues because susceptible learners shift more time from learning to production.

increase. The drop in output to the trough is less severe than in the economy without the lockdown. Therefore, the lockdown delays the peak of infections, flattens the infection curve, and reduces the fall in output caused by infections.

However, the welfare effect of the lockdown is small. In fact, the lockdown generates a loss of 0.25% in the permanent welfare flow. The lockdown generates a loss because it restricts the time input without reducing the accumulated deaths – it merely delays these deaths. This welfare result should be interpreted with great caution because the model abstracts from some potential benefits of a lockdown. For example, a lockdown may improve the ability to test and trace infections, and it reduces congestion in hospitals and medical facilities. Moreover, if a vaccine is expected to be developed, a lockdown may improve welfare significantly by reducing infections and deaths.

6. Conclusion

The Covid-19 pandemic is devastating to the world. Many lives have been lost to infections. Economic activities have ground to a halt as many countries have locked down. Although the number of infections worldwide has not declined substantially, countries have started to loosen the lockdown restrictions. The inescapable reality is that there is a tradeoff between containing the pandemic and maintaining economic activities. In this paper I have examined the tradeoff from the perspective of knowledge diffusion in the time of a pandemic. Learners contact teachers randomly to acquire the frontier knowledge that increases productivity. When one of the two individuals in a meeting carries an infectious pathogen while the other is susceptible, the former can transmit the pathogen to the latter. A learner chooses the learning intensity by taking into account the infection risk and the expected gain from learning. Calibrating the pathogen to Covid-19, the model shows that a susceptible learner temporarily curtails the learning intensity when the pathogen arrives, but soon increases the learning intensity above the level without the pathogen. As a result, knowledge diffusion slows down temporarily but soon increases in pace. Aggregate output falls by 13% and recovers in less than one year, and deaths reach 1.2% of the population.

A society is willing to give up 1.39% of permanent output or, equivalently, 31.6% of the first year's output to identify infectious learners and bar them from the learning process. Moreover, the transmission of the pathogen feeds on knowledge diffusion. If the economy starts in the steady state that requires less diffusion in knowledge, deaths caused by the same pathogen are 0.66% of the population, output falls by 4.4%, and the welfare gain from controlling infections is 12.2% of the first year's output.

Two extensions of the model in this paper are worthwhile pursuing. First, the model omits the possibility that a vaccine may be developed in the future. Although there are attempts in the literature to incorporate this possibility, they either do not model individuals' choices or assume that these choices are myopic.[17] But an expected vaccine will alter choices as individuals weigh the benefit of waiting for the vaccine against the cost of suppressed activities such as learning. As emphasized earlier, this forward-looking behavior should be central to answering the question of how long a society should maintain a lockdown to wait for a vaccine. Second, the model assumes that the pathogen does not mutate over time. Although this assumption may be reasonable for Covid-19, for now, it may not be so for a long time. As history shows, germs have always been with humans. As we learn and get smarter, germs also evolve. Allowing for the pathogen to mutate is especially important if one extends the current model to incorporate long-run growth. Such an extension may also enable one to tackle the age-old relationship between economic growth and the population that occupied the attention of Malthus (1798).

[17]Krasikov and Lamba (2020) construct an interesting model where a vaccine is expected to arrive with time-varying probabilities. However, they assume that individuals' choices depend only on the current and past variables but not on the expected values.

Appendix

A. Explanations and Proofs for Sections 2 and 3

The Explanations for (2.12)-(2.14):

The explanation for (2.12): This is the law of motion of n_{11t}, the group of susceptible teachers ($az = 11$). Susceptible learners join this group when they successfully learn and do not contract the pathogen in learning. This is the first term on the right-hand side of (2.12). The second term is the flow into the group by infected teachers who recover from the infections without immunity. Infections in the group generate a flow out of the group equal to $\psi\sigma^*_{bt}x_{0bt}n_{11t}$. The rate at which a teacher is contacted by an infectious learner is $\sigma^*_{bt}x_{0bt}$, as explained before for (2.5), and contracts the pathogen with the probability ψ. The last term in (2.12) is the measure of deaths in the group.

The explanation for (2.13): This is the law of motion of n_{1bt}, the group of infected teachers ($az = 1b$). The first three terms on the right-hand side of (2.13) are the flows into the group. They correspond, respectively, to infectious learners who succeed in learning, susceptible learners who succeed in learning and become infected in learning, and susceptible teachers who are infected. The outflows are deaths and the two types of recoveries.

The explanation for (2.14): This is the law of motion of n_{a0t}, the group of immune individuals with the knowledge index a. Infected individuals with knowledge a join the group when they acquire immunity, which generates an inflow μn_{abt}. Deaths generate an outflow from the group equal to $\delta_0 n_{a0t}$. If $a = 0$, individuals in the group are learners. Each has a meeting with a teacher at the rate σ^*_{0t} and learns with the probability β. This outflow from the group is equal to $\beta\sigma^*_{0t}n_{00t}$. If $a = 1$, individuals in the group are teachers, and the flow generated by individuals 00's learning is a flow into, rather than a flow out of the group. This change in the sign is captured by $1 - 2I_{a=0} = 1$ in this case.

Proof of Proposition 3.1

Proof. Examine equation (3.3) first. The right-hand of the equation is a function of only D_{0t} and the disease status. By the envelope condition on the optimal choice of ℓ_{0t},

the derivative of the right-hand side of (3.3) is equal to $r + \delta_0 + \beta\sigma\left(\ell_0\left(D_{0t}\right)\right) > 0$. Thus, (3.3) is an ordinary differential equation of D_{0t} with a strictly positive eigenvalue. Because D_{0t} is a jump variable, the solution for D_{0t} under rational expectations is that D_{0t} jumps immediately to the long-run level D_0^* that solves (3.4). Because the right-hand side of (3.4) is strictly increasing in D_0^*, as just calculated, the solution for D_0^* is unique if it exists. It is clear that the right-hand side of (3.4) is negative at $D_0^* = 0$, and positive at $D_0^* = \infty$. Thus, a unique solution for D_0^* exists indeed. Then, (2.8) implies $\ell_{0t}^* = \ell_0\left(D_0^*\right)$ for all t.

Setting $z = 0$ in (2.9), the equation becomes an ordinary differential equation of only v_{10t} with a positive eigenvalue. By a similar argument to the above, $v_{10t} = v_{10}^*$ for all t. Using this result and (2.9), I can calculate v_{10}^* as in (3.5). Then, the definition of Δ_0 and the result $\Delta_0^* = D_0^*/\beta$ imply $v_{00t} = v_{00}^* = v_{10}^* - D_0^*/\beta$ for all t. **QED**

Proof of Proposition 3.2

Proof. If $\rho = 0$, (3.6) becomes

$$\frac{\mathrm{d}}{\mathrm{d}t}D_{bt} = \left(r + \delta_b\right)D_{bt} - \mu\left(D_0^* - D_{bt}\right) - \beta\varepsilon_b L\left(1 - \gamma^{-1}\right) + \beta\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_{bt}\right]_{\ell=\ell_b(D_{bt})}$$

A similar proof to that of Proposition 3.1 applies here to D_b. That is, D_{bt} jumps immediately to the long-run level D_b^* that solves (3.7), and $\ell_{bt}^* = \ell_b\left(D_b^*\right)$ for all t. Similarly, if $\rho = 0$, then (2.9) with $z = b$ implies that $v_{1bt} = v_{1b}^*$, where v_{1b}^* is given in (3.8). This implies $v_{0bt} = v_{0b}^* = v_{1b}^* - D_b^*/\beta$ for all t.

To prove $v_{10}^* > v_{1b}^*$, I use the expressions for v_{10}^* in (3.5) and v_{1b}^* in (3.8) to derive:

$$\left(r + \delta_0 + \mu\right)\left(v_{10}^* - v_{1b}^*\right) = \left(\varepsilon_0 - \varepsilon_b\right)L + \left(\delta_b - \delta_0\right)v_{1b}^* > 0. \tag{A.1}$$

The inequality follows from $\varepsilon_0 > \varepsilon_b$, $\delta_b > \delta_0$ and $v_{1b}^* \geq 0$. By (3.5) and (3.8), $v_{00}^* > v_{0b}^*$ if and only if $D_0^* < D_b^* + \beta\left(v_{10}^* - v_{1b}^*\right)$. Substituting $\left(v_{10}^* - v_{1b}^*\right)$, this condition becomes

$$D_0^* < D_b^* + \frac{\beta\left[\left(\varepsilon_0 - \varepsilon_b\right)L + \left(\delta_b - \delta_0\right)v_{1b}^*\right]}{r + \delta_0 + \mu} \equiv D_a.$$

To prove that this inequality holds, subtract (3.4) and (3.7) to obtain:

$$\begin{aligned} 0 = \; & \left(r + \delta_0 + \mu\right)\left(D_0^* - D_b^*\right) - \left(\delta_b - \delta_0\right)D_b^* - \beta\left(\varepsilon_0 - \varepsilon_b\right)L\left(1 - \gamma^{-1}\right) \\ & + \beta\max_{\ell\in[0,\varepsilon_0 L]}\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_0^*\right] - \beta\max_{\ell\in[0,\varepsilon_b L]}\left[-\ell\gamma^{-1} + \sigma\left(\ell\right)D_b^*\right]. \end{aligned} \tag{A.2}$$

Temporarily denote the right-hand side of this equation as $f\left(D_0^*, D_b^*\right)$. Because f inherits the dependence of the right-hand of (3.4) on D_0^*, f is increasing in D_0^*. Then, $D_0^* < D_a$ if and only if $f\left(D_a, D_b^*\right) > 0$. Compute:

$$\begin{aligned}\tfrac{1}{\beta} f\left(D_a, D_b^*\right) = \; & \left(\delta_b - \delta_0\right) v_{0b}^* + \left(\varepsilon_0 - \varepsilon_b\right) L\gamma^{-1} \\ & + \max_{\ell \in [0, \varepsilon_0 L]} \left[-\ell\gamma^{-1} + \sigma\left(\ell\right) D_a\right] - \max_{\ell \in [0, \varepsilon_b L]} \left[-\ell\gamma^{-1} + \sigma\left(\ell\right) D_b^*\right].\end{aligned}$$

I have used the fact $v_{1b}^* = v_{0b}^* + D_b^*$. By the envelope condition of the first maximization problem, the maximum is (weakly) increasing in D_a. Because $D_a > D_b^*$, the maximum is (weakly) greater than the maximum achieved by the second maximization problem. Thus,

$$\frac{1}{\beta} f\left(D_a, D_b^*\right) \geq \left(\delta_b - \delta_0\right) v_{0b}^* + \left(\varepsilon_0 - \varepsilon_b\right) L\gamma^{-1} > 0.$$

This proves $D_0^* < D_a$ and, hence, $v_{00}^* > v_{0b}^*$.

To prove $D_0^* > D_b^*$, suppose $D_0^* \leq D_b^*$ to the contrary. By the envelope condition, the maximum of the first maximization problem in (A.2) cannot exceed that of the second maximization. In this case,

$$\begin{aligned} f\left(D_0^*, D_b^*\right) & \leq \left(r + \delta_0 + \mu\right)\left(D_0^* - D_b^*\right) - \left(\delta_b - \delta_0\right) D_b^* - \beta\left(\varepsilon_0 - \varepsilon_b\right) L\left(1 - \gamma^{-1}\right) \\ & \leq -\left(\delta_b - \delta_0\right) D_b^* - \beta\left(\varepsilon_0 - \varepsilon_b\right) L\left(1 - \gamma^{-1}\right) < 0.\end{aligned}$$

This contradicts the fact $f\left(D_0^*, D_b^*\right) = 0$.

Continue to assume $\rho = 0$. The above proof shows that for $z \in \{b, 0\}$, the optimal learning contact rate of a learner $0z$ is $\sigma\left(\ell_z\left(D_z\right)\right)$, where $\ell_z\left(D\right)$ solves $\max_{\ell \in [0, \varepsilon_z L]} \left[-\ell\gamma^{-1} + \sigma\left(\ell\right) D\right]$. The objective function is strictly supermodular in (ℓ, D), and the domain is strictly wider if ε_z is larger. Since $D_0^* > D_b^*$ and the solution to the maximization problem is unique, these features imply $\ell_0\left(D_0^*\right) \geq \ell_b\left(D_b^*\right)$ (see Topkis, 1998). The inequality is strict if both ℓ_0 and ℓ_b are interior or both are at the upper corner. **QED**

B. Analysis of the Planner's Problem

The planner's problem is specified in section 5.1. Let w_{azt} be the current-value multiplier of the law of motion of n_{azt}. This is the social marginal value of an individual az at time

t. The current value Hamiltonian of the planner's problem is:

$$\begin{aligned} H = & \sum\nolimits_{z=b,1,0} \left[n_{0zt} \left(\varepsilon_z L - \ell_{zt} \right) \gamma^{-1} + n_{1zt} \varepsilon_z L \right] - \zeta \left(\delta_b - \delta_1 \right) M_{bt} \\ & + w_{00t} \left[\mu n_{0bt} - \delta_0 n_{00t} - \beta \sigma_{0t} n_{00t} \right] + w_{10t} \left[\mu n_{1bt} - \delta_0 n_{10t} + \beta \sigma_{0t} n_{00t} \right] \\ & + w_{01t} \left[\frac{\bar{\delta}_t M_{1t}}{M_{bt}+M_{1t}} + \rho n_{0bt} - \left[\beta + (1-\beta) \psi x_{1bt} \right] \sigma_{1t} n_{01t} - \delta_1 n_{01t} \right] \\ & + w_{11t} \left[\left(1 - \psi x_{1bt} \right) \beta \sigma_{1t} n_{01t} + \rho n_{1bt} - \psi \sigma_{bt} x_{0bt} n_{11t} - \delta_1 n_{11t} \right] \\ & + w_{0bt} \left[\frac{\bar{\delta}_t M_{bt}}{M_{bt}+M_{1t}} + \psi x_{1bt} \left(1 - \beta \right) \sigma_{1t} n_{01t} - \beta \sigma_{bt} n_{0bt} - \left(\mu + \rho + \delta_b \right) n_{0bt} \right] \\ & + w_{1bt} \left[\beta \sigma_{bt} n_{0bt} + \psi \beta \sigma_{1t} x_{1bt} n_{01t} + \psi \sigma_{bt} x_{0bt} n_{11t} - \left(\mu + \rho + \delta_b \right) n_{1bt} \right] \end{aligned}$$

Note that $\bar{\delta}$, M_{zt} and x_{azt} depend on the n'es. In particular, $\bar{\delta}_t = \sum_{z=1,b,0} \delta_z M_{zt}$ and $x_{azt} = \frac{n_{azt}}{N_{1t}}$. The socially efficient choice of ℓ_{zt} satisfies the following complementary slackness, (5.3). In the condition, Q_{zt} is

$$\begin{aligned} Q_{0t} = & \left(w_{10t} - w_{00t} \right) \beta \\ Q_{1t} = & \left(w_{11t} - w_{01t} \right) \beta - \psi x_{1bt} \left[\left(w_{11t} - w_{1bt} \right) \beta + \left(w_{01t} - w_{0bt} \right) \left(1 - \beta \right) \right] \\ Q_{bt} = & \left(w_{1bt} - w_{0bt} \right) \beta - \psi x_{11t} \left[w_{11t} - w_{1bt} \right] \end{aligned}$$

The socially efficient measure of learners $0z$, n_{0zt}, satisfies the following conditions:

$$\begin{aligned} & \left(r + \delta_z \right) w_{0zt} \\ = & \frac{\mathrm{d}}{\mathrm{d}t} w_{0zt} + I_{z=b} \left[\mu \left(w_{00t} - w_{0bt} \right) + \rho \left(w_{01t} - w_{0bt} \right) \right] + \left(\varepsilon_z L - \ell_{zt} \right) \gamma^{-1} \\ & + \sigma_{zt} Q_{zt} + \delta_z \tfrac{w_{01t} M_{1t} + w_{0bt} M_{bt}}{M_{bt}+M_{1t}} + I_{z \neq 0} \left(w_{01t} - w_{0bt} \right) B_{zt} - I_{z=b} \zeta \left(\delta_b - \delta_1 \right). \end{aligned}$$

The above equation resembles the Bellman equation for v_{0zt} in (2.6). The social marginal value of an individual, w, replaces the private value, v, and the social gain from learning, Q, replaces the private gains, D. However, the last three terms in the above equation do not have counterparts in the Bellman equation for v_{0zt}. They all arise because the planner takes into account the dependence of deaths and births on the current distribution of individuals. The weighted average of (M_{1t}, M_{bt}) multiplied by δ_z is the effect of n_{0zt} on the measure of deaths, $\bar{\delta}_t$, and, hence, on the measure of births. The term with $I_{z \neq 0}$ is the effect of n_{0zt} on the composition of births, where

$$B_{zt} \equiv \bar{\delta}_t \frac{\partial}{\partial n_{0zt}} \left[\frac{M_{1t}}{M_{bt} + M_{1t}} \right] \text{ for } z = 1, b.$$

This term exists only if $z \neq 0$, because there are no immune individuals among newborns. The last term with $I_{z=b}$ in the equation for w_{0zt} is the effect of n_{0zt} on the additional deaths caused by infections, which exists only for $z = b$. Note that the planner does incorporate

the death rate as part of the effective discounting, as shown by the left-hand side of the equation for w_{0zt}, because deaths affect the laws of motion of the distribution.

Similarly, the socially efficient measure of teachers $1z$, n_{1zt}, satisfies the following condition:

$$\begin{array}{rl}(r+\delta_z)\,w_{1zt} = & \frac{\mathrm{d}}{\mathrm{d}t}w_{1zt} + I_{z=b}\left[\mu\left(w_{10t}-w_{1bt}\right)+\rho\left(w_{11t}-w_{1bt}\right)\right]+\varepsilon_z L \\ & +\delta_z\frac{w_{01t}M_{1t}+w_{0bt}M_{bt}}{M_{bt}+M_{1t}} + I_{z\neq 0}\left(w_{01t}-w_{0bt}\right)B_{zt} - I_{z=b}\zeta\left(\delta_b-\delta_1\right)+K_t.\end{array}$$

This equation resembles the Bellman equation of v_{1zt} in (2.9), with the social gains and marginal values replacing the private ones. Similar to the equation for w_{0zt}, the planner incorporates the effect of n_{1zt} on the measure of births, $\bar{\delta}_t$, the composition of newborns, and the additional deaths from infections. The term K_t is the effect of n_{1zt} on the social marginal value through the total measure of teachers, N_{1t}, which appears in the laws of motion through (x_{0bt}, x_{1bt}). Precisely,

$$K_t = \psi\sigma_{bt}x_{0bt}x_{11t}\left[w_{11t}-w_{1bt}\right]+\psi\sigma_{1t}x_{1bt}x_{01t}\left[\left(w_{11t}-w_{1bt}\right)\beta+\left(w_{01t}-w_{0bt}\right)\left(1-\beta\right)\right].$$

C. The Calibration and Computation

The calibration procedure:

The frequency for the calibration is weekly. The amount of time available to an individual is $L = 1$. The target on the annual interest rate implies $e^{52r} = 1.05$, which solves the weekly interest rate r. The parameters (η, γ, b) are set to values as discussed in the main text. The death rate of an immune individual, δ_0, is such that the expected length of the working life is 65 years. That is, $1/\delta_0 = 65 \times 52$. The death rate of a susceptible individual is $\delta_1 = \delta_0$.

The pathogen is calibrated to Covid-19. The expected length to a recovery from an infection is 4 weeks; i.e., $\frac{1}{\mu+\rho} = 4$. Conditional on a recovery in this mean duration, the probability that the recovery comes with immunity is 0.8. That is, $\frac{1-e^{-\mu/(\mu+\rho)}}{1-e^{-1}} = 0.8$. These two conditions solve ρ and μ. The death probability in the average duration of an infection is $1-e^{-\delta_b/(\mu+\rho)}$. The death probability in the same length of time for an individual without the infection is $1-e^{-\delta_1/(\mu+\rho)}$. The additional death probability caused by an infection is

$e^{-\delta_1/(\mu+\rho)} - e^{-\delta_b/(\mu+\rho)}$. Setting this to 0.06 solves δ_b. The expected length of an infection is the incubation period, 2 weeks. The probability ψ is equal to the ratio of the incubation period to the expected length to a recovery; i.e., $\psi = 2/4$. The probability of successfully learning the frontier knowledge in a meeting, β, and the scale parameter in the contact rate function, σ_c, jointly meet two targets. One is that the expected length of time needed to acquire the frontier knowledge if learning full time is 5 years. That is, $\frac{1}{\beta\sigma_c} = 5 \times 52$. The other is that the expected length of time it takes to be infected if a learner makes contacts full time is 4 days. That is, $\frac{1}{\psi\sigma_c} = \frac{4}{7}$. Since ψ is already determined, the two equations solve (β, σ_c). The target on the number of days needed to be infected also generates the effective reproductive number in the first 10 weeks to be 3 on average.

The Computation Procedure:

The computation procedure consists of Steps 1-5 below. The inputs into the computation are the parameter values, the initial distribution of individuals, and the experiment to be conducted. The procedure handles the general case where the path of the time endowment is $\{L_t\}_{t\geq 0}$ and the path of the contact efficiency is $\{\sigma_{ct}\}_{t\geq 0}$. For example, a lockdown affects both paths. In the absence of a lockdown, $L_t = 1$ and $\sigma_{ct} = \sigma_c$ for all t.

Step 1. Compute the policy function of the optimal choice, $\ell^*_{zt} = \ell_z(D_{zt})$:

$$\sigma_{ct}\eta\left(\ell^*_{zt}\right)^{n-1}\gamma\hat{D}_{zt} \geq 1 \text{ and } \ell^*_{zt} \leq \varepsilon_z L_t.$$

Denote the implied contact rate as σ^*_{zt} and the learning intensity as $s^*_{zt} = \ell^*_{zt}/\varepsilon_z$.

Step 2. Construct the Chebyshev grid of time t, $\{t_1, t_2, ..., t_M\}$, and the Chebyshev basis that contains N ($\leq M-1$) Chebyshev polynomials. These will be used for approximating functions of t.

Step 3. Solve the functions v_{00t} and v_{10t} by iterating on the Bellman equations. These equations generalize (2.6) and (2.9) for $z = 0$ to allow for time variations in (L_t, σ_{ct}):

$$(r+\delta_0)\,v_{00t} = RHS_{00t}, \quad (r+\delta_0)\,v_{10t} = RHS_{10t},$$

where

$$\begin{aligned} RHS_{00t} &= \tfrac{\mathrm{d}}{\mathrm{d}t}v_{00t} + \varepsilon_0 L_t\gamma^{-1} + \left[-\ell^*_{0t}\gamma^{-1} + \sigma\left(\ell^*_{0t}\right)D_{0t}\right] \\ RHS_{10t} &= \tfrac{\mathrm{d}}{\mathrm{d}t}v_{10t} + \varepsilon_0 L_t \end{aligned}$$

For the iteration, let $\mathrm{d}t > 0$ be a length of time and rewrite the Bellman equations as

$$v_{a0t} = \frac{v_{a0t} + RHS_{a0t}\mathrm{d}t}{1 + (r + \delta_0)\,\mathrm{d}t},\ a = 0, 1. \tag{C.1}$$

Note that these are equivalent to the original Bellman equations, not an approximation. The following sub-loop (v1)-(v4) performs the iteration on (C.1):

(v1) Start with an initial guess of (v_{00t}, v_{10t}) as functions of t. Construct the Chebyshev coefficients of these functions on the Chebyshev basis.

(v2) Compute D_{0t}, and evaluate ℓ^*_{0t} and σ^*_{0t}.

(v3) Use this initial guess to compute the right-hand side of (C.1). In particular, approximate the derivative $\frac{\mathrm{d}}{\mathrm{d}t}v_{a0t}$ in RHS_{a0t} by the forward difference:

$$\frac{\mathrm{d}}{\mathrm{d}t}v_{a0t} \approx \frac{v_{a0(t+d)} - v_{a0t}}{d},\ d > 0.$$

Note that $t + d$ is usually off the grid of t. Use Chebyshev projection to compute the value $v_{a0(t+d)}$.

(v4) Use (C.1) and the result in (v3) to update (v_{00t}, v_{10t}). Return to (v1) and repeat the procedure until convergence.

Step 4. Solve the path of the distribution $n_t = (n_{0bt}, n_{1bt}, n_{01t}, n_{11t}, n_{00t}, n_{10t})$ and the value functions $(v_{abt}, v_{a1t})_{a=0,1}$. The iteration follows steps (n1)-(n5):

(n1) Make an initial guess of the path of n_t, with n_0 being given. Ensure that each element of n_t is bounded in $[0, 1]$ and the elements sum up to 1.

(n2) Solve the value functions $(v_{abt}, v_{a1t})_{a=0,1}$ by using a sub-loop similar to (v1)-(v4) above. In (v2), the expected gains are D_{bt} and D_{1t}, and the policy functions to be evaluated are $(\ell^*_{bt}, \ell^*_{1t})$. The implied contact rates are $(\sigma^*_{bt}, \sigma^*_{1t})$. Note that these gains and policy functions depend on the distribution, which explains why the value functions $(v_{abt}, v_{a1t})_{a=0,1}$ have to be computed inside the iteration of n_t rather than outside. The Bellman equations used in a step similar to (v3) are

$$v_{abt} = \frac{v_{abt} + RHS_{abt}\mathrm{d}t}{1 + (r + \delta_b)\,\mathrm{d}t}, \quad v_{a1t} = \frac{v_{a1t} + RHS_{a1t}\mathrm{d}t}{1 + (r + \delta_1)\,\mathrm{d}t}, \quad a = 0, 1$$

where

$$
\begin{aligned}
RHS_{0bt} &= \tfrac{\mathrm{d}}{\mathrm{d}t} v_{0bt} + \mu \left(v_{00} - v_{0bt} \right) + \rho \left(v_{01t} - v_{0bt} \right) + \left(\varepsilon_b L_t - \ell_{bt}^* \right) \gamma^{-1} + \sigma \left(\ell_{bt}^* \right) D_{bt} \\
RHS_{1bt} &= \tfrac{\mathrm{d}}{\mathrm{d}t} v_{1bt} + \mu \left(v_{10} - v_{1bt} \right) + \rho \left(v_{11t} - v_{1bt} \right) + \varepsilon_b L_t \\
RHS_{01t} &= \tfrac{\mathrm{d}}{\mathrm{d}t} v_{01t} + \left(\varepsilon_1 L_t - \ell_{1t}^* \right) \gamma^{-1} + \sigma \left(\ell_{1t}^* \right) D_{1t} \\
RHS_{11t} &= \tfrac{\mathrm{d}}{\mathrm{d}t} v_{11t} + \varepsilon_1 L_t - \psi \sigma_b^* x_{0bt} \left[v_{11t} - v_{1b} \right].
\end{aligned}
$$

Note that the terms with μ and ρ in the equations for (v_{0bt}, v_{1bt}) are kept on the right-hand side instead of being added to the effective discount rate. This makes the precision and the convergence speeds of different value functions comparable.

(n3) Given the initial distribution n_0, solve the differential equations (2.10)-(2.14) for n_t. In these equations, the contact rates are those generated by the optimal choices in the last round of iteration in (n2) above.

(n4) Update n_t by the solution in (n3). Return to (n1) and repeat until convergence.

Step 5. Compute other statistics and equilibrium variables, such as the expected reproductive number and aggregate output. In particular, the expected reproductive numbers, R_{1bt} and R_{0bt}, obey the differential equations (2.15) and (2.16). These are forward-looking equations similar to the Bellman equations. They are solved with the same method as that for the value functions, described by (v1)-(v4) above.

For the planner's problem, the computation is similar, with the corresponding value functions and Bellman equations. The modification is that the social marginal value of immune individuals, w_{00t} and w_{10t}, can depend on the distribution of individuals. As a result, they need to be computed together with other functions inside the loop for the distribution, as in Step 4 above.

References

[1] Acemoglu, D. Chernozhukov, V., Werning, I. and M.D. Whinston, 2020, "A Multi-Risk SIR Model with Optimally Targeted Lockdown," NBER working paper 27102.

[2] Alvarez, F., Argente, D. and F. Lippi, 2020, "A Simple Planning Problem for COVID-19 Lockdown and Smart Tracing," working paper, Pennsylvania State University.

[3] Anderson, R.M. and R.M. May, 1979, "Population Biology of Infectious Diseases: Part I," Nature 280, 361-367.

[4] Atkeson, A., 2020, "What Will Be the Economic Impact of Covid-19 in the US? Rough Estimates of Disease Scenarios," NBER working paper 26867.

[5] Ben-Porath, Y., 1967, "The Production of Human Capital and the Life Cycle of Earnings," Journal of Political Economy 75, 352-365.

[6] Berger, D.W., Herkenhoff, K.F. and S. Mongey, 2020, "An SEIR Infectious Disease Model with Testing and Conditional Quarantine," NBER Working paper 26901.

[7] Brotherhood, L., Kircher, P., Santos, C. and M. Tertilt, 2020, "An Economic Model of the Covid-19 Epidemic: The Importance of Testing and Age-Specific Policies," working paper, European University Institute.

[8] Diamond, J., 1999, Guns, Germs, and Steel. New York, W.W. Norton & Company.

[9] Dworsky, M., Seabury, S.A., Neuhauser, F.W., Kharel, U., and R. Euller, 2016, "Benefits and Earnings Losses for Permanently Disabled Workers in California: Trends Through the Great Recession and Impacts of Recent Reforms," RAND Center for Health and Safety in the Workplace.

[10] Eichenbaum, M., Rebelo, S. and M. Trabandt, 2020, "The Macroeconomics of Epidemics," NBER Working paper 26882.

[11] Farboodi, M., Jarosch, G. and R. Shimer, 2020, "Internal and External Effects of Social Distancing in a Pandemic," working paper, University of Chicago.

[12] Greenwood, J., Kircher, P, Santos, C. and M. Tertilt, 2019, "An Equilibrium Model of the African HIV/AIDS Epidemic," Econometrica 87, 1081-1113.

[13] Hall, R.E., Jones, C. and P.J. Klenow, 2020, "Trading off Consumption and COVID-19 Deaths," working paper, Stanford University.

[14] Kapicka, M. and P. Rupert, 2020, "Labor Markets during Pandemics," working paper, University of California at Santa Barbara.

[15] Keppo, J., Quercioli, E., Kudlyak, M., Smith, L. and A.Wilson, 2020. "The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics," working paper, Federal Reserve Bank of San Francisco.

[16] Kermack, W.O. and A.G. McKendrick, 1927, "A Contribution to the Mathematical Theory of Epidemics," Proceedings of the Royal Society of London A 115, 700-721.

[17] Kirkcaldy, R.D., King, B.A. and J.T. Brooks, 2020, "COVID-19 and Postinfection Immunity: Limited Evidence, Many Remaining Questions," JAMA, doi:10.1001/jama.2020.7869

[18] Krasikov, I. and R. Lamba, 2020, "Behavioral Epidemiology: An Economic Model to Evaluate Optimal Policy in the Midst of a Pandemic," working paper, Pennsylvania State University.

[19] Light, A., 2005, "Job Mobility and Wage Growth: Evidence from the NLSY79," Monthly Labor Review, the U.S. Bureau of Labor Statistics.

[20] Linka, K., Peirlinck, M. and E. Kuhl, 2020, "The Reproduction Number of COVID-19 and Its Correlation with Public Health Interventions," medRxiv, https://doi.org/10.1101/2020.05.01.20088047

[21] Lucas, R.E., Jr. and B. Moll, 2014, "Knowledge Growth and the Allocation of Time," Journal of Political Economy 122, 1-51.

[22] Malthus, T.R., 1798, An Essay on the Principle of Population.

[23] Topkis, D.M., 1998, Supermodularity and Complementarity. Princeton, NJ: Princeton University Press.

[24] Toxvaerd, F., 2020, "Social Distancing with Asymptomatic Infection: Beliefs, Fatalism and Testing," working paper, University of Cambridge.

[25] U.S. Environmental Protection Agency, 2018, "Mortality Risk Valuation," https://www.epa.gov/environmental-economics/mortality-risk-valuation#means