



ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS
WORKING PAPER SERIES

Optimal Management of an Epidemic: Lockdown, Vaccine and Value of Life

Authors	Carlos Garriga, Rodolfo E. Manuelli, and Siddhartha Sanghi
Working Paper Number	2020-046B
Revision Date	May 2022
Citable Link	https://doi.org/10.20955/wp.2020.046
Suggested Citation	Garriga, C., Manuelli, R.E., Sanghi, S., 2022; Optimal Management of an Epidemic: Lockdown, Vaccine and Value of Life, Federal Reserve Bank of St. Louis Working Paper 2020-046. URL https://doi.org/10.20955/wp.2020.046

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

Optimal Management of an Epidemic: Lockdown, Vaccine and Value of Life*

Carlos Garriga[†] Rody Manuelli[‡] Siddhartha Sanghi[§]

First Version: April 27, 2020

This Version: February 24, 2022

[\[Most Updated Version Here\]](#)

Abstract

This paper analyzes the optimal management of a pandemic (stay-at-home and vaccination policies) in a dynamic model. The optimal lockdown policies respond to the spread of the virus with significant restrictions to employment, followed by partial loosening before the peak of the epidemic. Upon the availability of a vaccine, the optimal vaccination policy has an almost bang-bang property, despite the loss of immunity of the vaccinated: vaccinate at the highest possible rate, and then rapidly converge to the steady state. The model illustrates interesting trade-offs as it implies that lower hospital capacity requires flattening the infection curve and hence a more stringent lockdown, but lower vaccination possibilities (both the likelihood of a vaccine and the vaccination rate) push the optimal lockdown policy in the opposite direction, even before the arrival of vaccine. The model implies that the “dollar” value of a vaccine decreases rapidly as time passes with the reinfection rate being an important determinant of the monetary value. The value that society assigns to averting deaths is a major driver of the optimal policy. The sensitivity analysis shows that even for reasonable bounds of the economic and epidemiological parameters, the timing and the magnitude of the optimal policy varies substantially.

*The views expressed in this paper do not necessarily reflect the position of the Federal Reserve Bank of St. Louis. We thank Emilio Espino, Miguel Faria-e-Castro, Piero Gottardi, Fernando Navajas, Andy Neumeyer, David Slusky and Chris Waller as well as participants in several seminars for their comments. Manuelli and Sanghi want to thank the Weidenbaum Center for financial support. Sanghi has also received support from the Koch Center for Family Business.

[†]Federal Reserve Bank of St. Louis

[‡]Washington University in St. Louis

[§]Federal Reserve Bank of St. Louis

1 Introduction

How does the optimal response to an epidemic—both in terms of the nature of “stay-at-home” policies that restrict employment as well as the intensity of vaccination efforts—depend on the features of the economy and the epidemiological parameters? To make progress in understanding the answer to this question, we study the problem faced by a planner that can impose restrictions on employment and, when a vaccine becomes available, has to decide on the (costly) intensity at which the population can be vaccinated.

We study a standard continuous time infinite horizon model. We assume that individual preferences depend on individual consumption, and that social preferences take into account the utility loss associated with deaths. In addition, we model the impact of a virus using a standard SIRS model, and we constrain the policy space taking into account limitations imposed by the existing public health infrastructure. Another important dimension of the model is that a fraction of individuals that have been vaccinated lose immunity. That requires additional rounds of vaccination for individuals previously vaccinated.

We assume that at the beginning of the epidemic—what we label Phase I—the only policy available to the planner is a stylized version of a “stay-at-home” policy that, simultaneously, restricts employment and lowers the rate of transmission of the epidemic. Phase I ends when a vaccine becomes available and the economy enters Phase II. We view the arrival of a vaccine as a random event and take the probability distribution as exogenous. At this point, the planner has a second tool to control the epidemic: the speed at which the population can be vaccinated, which we also view as requiring resources. This less-than-instantaneous ability to vaccinate the population is a novel feature of our model and one that has significant effects on the optimal policy, even before the vaccine becomes available.

On the theoretical side, we show that the model has a steady state. In the case where the reinfection rate is low (essentially the case in which the economy is dealing with an epidemic and not an endemic problem), we show that along a path in which a vaccine or a treatment never becomes available (Phase I)—although optimal policies take into account that the probability is positive—the epidemiological variables converge to the *same steady state* as those of another economy that has access to a vaccine/treatment. This implies that the economic value of a vaccine decreases over time. To the extent that the private value of a vaccine moves with the social value, the model predicts that fewer resources will be allocated by the private sector to finding a vaccine as the epidemic progresses.

We study a quantitative version of the model. We consider a variety of scenarios to capture the uncertainty associated with the true value of epidemiological parameters, the effective availability of health care resources in the case of an epidemic, and the differential case fatality rates associated with situations in which hospital capacity is exceeded. Our findings about optimal policies (both employment and vaccination) are very sensitive to assumptions about the appropriate value of the relevant parameters (epidemiological as well as public health infrastructure).

We find that in the initial phase (Phase I) it is generally optimal to impose large restrictions on employment in the initial weeks of the epidemic. Depending on epidemiological parameters, the restrictions are slowly removed. In all cases employment is increased *before* the epidemic peaks. We also find that the lower the probability of a vaccine or the lower the speed at which the population can be vaccinated, the *less*-restrictive are the employment policies. The reason for this is standard: If there is no tool to fight the virus there is no need to flatten the curve as the same number of individuals will perish. The only force that counterbalances this result is the assumption that the case fatality rate increases when hospital capacity is exceeded. In this case, flattening the curve can result in fewer deaths. Thus, the optimal policy has to balance these different factors that have opposing effects.

The details of how a vaccine interacts with other policies is novel and interesting as arrival of a vaccine (Phase II) *does not* imply—depending on the state of the epidemic—that all restrictions on employment should be lifted. Moreover, availability of a vaccine may result in an *increase* in the spread of the epidemic. This last somewhat counterintuitive result can be easily explained: Availability of a vaccine increases the rate at which the susceptible population shrinks and this reduces the future contagion rate. This implies that the cost of the epidemic in terms of future deaths and consumption decreases (less future contagion) and that, consequently, the marginal cost in terms of current output should decrease as well. This last step requires a liberalization (more contact among individuals) that, in turn, pushes up the contagion rate. Our results imply that for a developed country like the U.S., it is in general optimal to vaccinate at the highest possible rate (institutionally determined) when the vaccine first becomes available. The optimal vaccination policy rapidly converges to its long-run value (which is non-zero in the case of reinfections). Thus, the optimal vaccination policy has almost a bang-bang quality.

The model implies that for a large range of plausible scenarios the social value of a vaccine decreases rapidly as time goes by. In our quantitative model (that allows for the endemic nature of the virus) the dollar value of a vaccine decreases by about 60% after one year. The specific dollar value depends critically on the reinfection rate and the social value of life.

As mentioned before, we study a fairly large range of plausible scenarios. Even though we find our quantitative results useful we are fully aware that their *quality is no better than the quality of the data that we use*. At this point there is *significant uncertainty* about many of the key parameters, both those corresponding to the economic model as well as those implicit in the epidemiological mode. Our findings suggest that the optimal policy is *very sensitive* to the specific parameterization. It is not clear to us what to conclude from this other than showing the importance of acquiring information (for example, adopting a large program of random testing). We also view our results as providing policymakers with a framework of reference to study worst-case scenarios.

Our work falls within the large and growing macro literature that emphasizes the trade-off between managing the epidemic by controlling the spread of the infection and economic outcomes. There is a large (and growing) number of papers that use optimal control techniques

to explore the management of an epidemic.¹ This paper is closest to the recent work of Alvarez, Argente, and Lippi (2020), Acemoglu et. al (2020) and Gonzalez-Eiras and Niepelt (2020). The major difference from the other planner models is that we take a different approach to modeling the effect of the availability of a vaccine, and this allows us to evaluate the consequences of different arrival times. This permits us to discuss how the optimal policy and the value of a vaccine depends on both the epidemiological variables as well as the implicit value of life. We also formalize the role of hospital capacity, explicitly allowing for reinfections due to loss of immunity, both of which are key determinants of optimal policy. Even though our formulation emphasizes a control approach, in appendix 4 we sketch a simple model that allows for a behavioral component and captures individual response to the risks associated with infection (see, for e.g., Bisin and Gottardi (2021) for a model with externalities).²

Section 2 presents the model and section 3 discusses some theoretical results. Section 4 presents our quantitative findings. Section 5 briefly discusses ongoing work on extensions, and section 6 offers some preliminary concluding comments.

2 Model

We study a standard continuous time macro model. We assume that there is one good that is produced exclusively with labor. There are two policy variables that we study. First, a type of “stay-at-home” limit on the utilization rate of the labor force, which has two impacts: It decreases output and, simultaneously, reduces the rate of transmission of a virus since fewer individuals are in contact with others. The second policy is the rate at which individuals can be vaccinated when a vaccine becomes available. This rate is also subject to an institutional constraint that captures both delays in producing a viable vaccine in large quantities (even after one has been discovered) and the logistical arrangements associated with mass vaccination.

We assume that there is a representative agent that cares about consumption. Social pref-

¹See Alvarez, Argente, and Lippi (2020), Gonzalez-Eiras and Niepelt (2020), Acemoglu, Chernozhukov, Werning, and Whinston (2020), Jones, Philippon, and Venkateswaran (2020).

²Acemoglu et al. study optimal lockdown for heterogenous agents and find that the optimal policy calls for different lockdown strategies for different individuals. Gonzalez-Eiras and Niepelt present a general model but they concentrate on special cases in order to find closed-form solutions. Appendix 4 has a partial list of the rapidly growing literature on the economic effects of COVID-19. Other papers explore the effects of different policies in the dynamics of the epidemic (i.e., social distancing, lockdown, testing, etc.) in SIR models developed by Kermack and McKendrick (1927). See for example, Adda (2016), Atkeson (2020a,2020b), Aum, Lee, and Shin (2020), Avery, Bossert, Clark, Ellison, and Ellison (2020), Azzimonti et al. (2020), Baqaee et al. (2020), Bassetto (2020), Berger, Herkenhoff, and Mongey (2020), Bodenstein et al. (2020), Chang and Velasco (2020), Droz and Tavares (2020), Eichenbaum, Rebelo, and Trabandt (2020a,b),Farboodi, Jarosch, and Shimer (2020), Fang, Wang, and Yang (2020), Glover, Heathcote, Krueger, and Ríos-Rull (2020), Hsiang et al. (2020), Neumeyer (2020), Piguillem and Shi (2020), Pindyck (2020), Shao (2020), Wang et al. (2020). In this paper, the labor market is fairly stylized, but this is something that has been addressed in other papers that abstract from the optimal policy. See, for example, Bick and Blandin (2020), Kapicka and Rupert (2020), Dingel and Neiman (2020), Kurman, Lale, and Ta (2020).

erences are simply individual preferences adjusted (downward) by the disutility cost of deaths. Thus, from society’s perspective there are two reasons to control an epidemic: the direct loss of output associated with lower labor force availability and, in our baseline, the additional disutility cost of deaths associated with the epidemic. The details of how we model this disutility are spelled out below.

We consider two phases that differ in the availability of a vaccine.

- *Phase I*: This is the period in which there is no vaccine available. The only available tool is a “stay-at-home” type of policy that reduces employment. We use a single variable to capture a variety of interventions that affect both the rate of transmission of the virus and the level of employment. We leave for future work the analysis of policies that are likely to vary in their impact, such as social distancing, age-related limitations and complete lockdown, among others.
- *Phase II*: We assume that the availability of a vaccine arrives at an exogenous rate and at that time the economy enters Phase II. The planner has, in addition to the lockdown policy, the ability to control (up to a maximum) the speed at which the population is vaccinated.³

2.1 The Economic Model

We assume that there is only one good that is produced linearly using labor. If the available labor force is denoted L and only a fraction $\phi \in [0, 1]$ is utilized in production, utility is $u(\phi wL - c_V(\mu(1 - V)))$, where the second term captures the cost in terms of output in vaccinating a population of size $1 - V$. This is the unvaccinated population. This term is operative only in Phase II when a vaccine is available.

Social preferences depend on the utility derived from consumption (we abstract away from leisure at this stage) and an additional term that captures the disutility associated with the loss of life. The instantaneous social payoff is

$$u(\phi wL - c_V(\mu(1 - V))) - \Delta(D^1, D^+).$$

We make standard assumptions about the utility function u . In the simple model, L equals the (fixed) labor force minus those infected individuals who have been identified as such. In general we assume that the function $\Delta(D^1, D^+)$ is increasing and convex. The variables (D^1, D^+) —more thoroughly described later—capture case fatality rates depending on whether the hospital capacity is reached.

Society’s preferences are then a function of consumption and deaths. Let T_η be the (random) time at which the economy transitions to Phase II (that is, when vaccination becomes available).

³Alvarez et al. (2020) assume that most of the population can be instantaneously vaccinated.

Formally, preferences are given by

$$U = E\left\{ \int_0^{T_\eta} e^{-\rho t} [u(\phi_t w L_t) - \Delta(D_t^1, D_t^+)] dt \right. \\ \left. + e^{-\rho T_\eta} \int_0^\infty e^{-\rho t} \left[u(\phi_{T_\eta+t} w L_{T_\eta+t} - c_V(\mu_{T_\eta+t} (S_{T_\eta+t} + (1-\zeta) I_{T_\eta+t}))) - \Delta(D_{T_\eta+t}^1, D_{T_\eta+t}^+) \right] dt \right\} \quad (1)$$

where the expectation is taken over the realization of T_η .

Our description of the model ignores any behavioral component, that is, formally, it ignores how individuals adjust their behavior—in our case the only relevant dimension is their labor supply decision—depending on the risk of infection. In the Appendix we sketch a version of the model that includes behavioral components and that, under a particular transfer policy, aggregates to the version that we use in the text.⁴

2.1.1 Special Case

The special case that we use for the quantitative exercise assumes that the instantaneous payoff is

$$N \ln(w\phi L - \underline{c}) - N\Delta(D^1, D^+).$$

In this formulation, \underline{c} is the minimum level of consumption and N is population size. We assume that the function $\Delta(D^1, D^+)$ has the following form:

$$\Delta(D_t^1, D_t^+) = M_0 [k_a D_t^1 + k_e D_t^+], \quad (2)$$

where

$$D_t^1 = \chi h \zeta \kappa I, \\ D_t^+ = (\chi^H - \chi) \kappa^H (h \zeta I - \bar{H})^+$$

This formulation has two elements. The cost to society of one additional death (in utility terms) is M_0 . If we assume that this is equal to the utility of the remaining lifetime T , and the value of an additional year is a multiple, v , of annual output, then we can approximate the utility loss associated with one death as

$$M_0 = \ln(vw - \underline{c}) \frac{1 - e^{-\rho T}}{\rho}.^5$$

We take D_t^1 as a measure of deaths that occur *before* hospital capacity is reached. To simplify, we view this flow as the product of the flow of people leaving the symptomatic infectious state, $\zeta \kappa I$, times the fraction hospitalized, h , times the fraction of those hospitalized that die, χ . D_t^+ has a similar interpretation, except that the death rate, $(\chi^H - \chi)$, is higher. We also allow

⁴We thank Piero Gottardi for encouraging us to think about the behavioral dimension.

⁵There are different approaches to identify and measure the valuation of life. Examples of some of the options in the context of the analysis of epidemics are Greenstone and Nigam (2020) and Hall, Jones and Klenow (2020).

for a shorter duration (κ^H). This term is only effective when the number of people hospitalized exceeds the hospital capacity. This is captured by the term $(h\zeta I - \bar{H})^+$.⁶ Total case fatalities, D_t , is simply the sum of D_t^1 and D_t^+ . Finally, the parameters k_a and k_e allow for the possibility that society values deaths that occur because hospital capacity is exceeded differently from those deaths that occur when hospitals have available beds.

2.2 The Epidemiological Model

Following the literature, we assume that the dynamics of an epidemic can be reasonably approximated by a version of the standard *SIR* model, augmented to allow for vaccination, making it *SIRV*.⁷

Let S be the number of susceptible individuals. I is the total number of infectious individuals. This includes both symptomatic and asymptomatic. We assume that only a certain fraction, ζ , is identified as infected/symptomatic. These individuals do not contribute to the labor supply, and we assume that they do not infect susceptible agents. Effectively, we assume that they are quarantined. The number of infected individuals who are asymptomatic is then $(1 - \zeta)I$. R the population of resistant individuals. In Phase II there exists another category: vaccinated individuals, which we denote by V .

Therefore, the potential labor force, L , is given by

$$L = S + (1 - \zeta)I + R + V. \quad (3)$$

Since we normalized the population to one, this is

$$L = 1 - \zeta I. \quad (4)$$

Then given a value of the stay-at-home policy ϕ , the fraction of susceptibles and infectious in the population is ϕS and $\phi(1 - \zeta)I$ respectively.⁸

⁶It is not obvious that the right approach is to posit that total rather than “excess” deaths should enter social preferences. For example, a large number of individuals die every year due to simple influenza. At the same time, there are relatively simple policies that could potentially avert many of those deaths (e.g., free vaccination and creating “vaccination stations” in convenient places, such as supermarkets and public transportation hubs, to reach a large fraction of the population, including those that do not have ready access to healthcare). We view the absence of those policies as a revealed preference type of argument against including all deaths in the baseline.

⁷One of the most widely cited epidemiological studies of the COVID-19 epidemic is the Imperial College model in Ferguson et.al. that uses the SIR model. Economic analyses of the COVID-19 epidemic from an economic point of view relying on the SIR model include Alvarez et al. (2020), Atkeson (2020), Fernandez-Villaverde and Jones (2020).

We are aware of the limitations of the model. See Korolev (2020) for example. An alternative forecasting model, the IHME model also appears to have serious limitations. See Marchant et al. (2020)

⁸In this setting, ϕ is a summary of the effects of a variety of different policies like lockdown, social distancing, school closure, mask wearing, travel restrictions and centralized quarantine. There is some evidence (see Chen

The model is then described by the following equations:

$$\begin{aligned}\dot{S} &= -\beta(\phi S)(\phi(1 - \zeta)I) - \mu S - \pi S + \gamma(1 - S) \\ \dot{S} &= -\beta\phi^2(1 - \zeta)SI - \mu S - \pi S + \gamma(1 - S).\end{aligned}\tag{5}$$

The first term is the standard matching function of the *SIRS* model, while the second term, μS , is the population that becomes resistant as a result of vaccination.⁹ The term πS is a measure of the flow of individuals that, even in the absence of contact, leave the *S* state because they become infected. This captures exogenous sources of infection. In our quantitative exercise we assume that π is such that the expected duration of the time in between infections is 40 years. The last term, $\gamma(1 - S)$, captures both the rate at which resistant individuals lose their immunity and the entrance of new susceptible individuals into the population.

The stock of infectious evolves according to

$$\dot{I} = \beta\phi^2(1 - \zeta)SI + \pi S - (\kappa + \mu + \gamma) I.\tag{6}$$

The last term includes both those who leave the infectious state because they either recover or die, and the fraction of the asymptomatic infectious population that is vaccinated. We assume—but there is mixed evidence on this—that vaccinated individuals cannot transmit the virus. It is easy enough to consider the case in which only a fraction effectively become vaccinated.

Finally the stock of vaccinated individuals satisfies

$$\dot{V} = \mu(1 - V) - \gamma V.\tag{7}$$

We assume that the population of candidates for vaccination includes those individuals who have not been vaccinated before and those that have lost their immunity. In the model, V is intended to capture individuals who have immunity. Thus, if \tilde{V} is the stock of individuals who have been vaccinated (and have not lost immunity) and if ζ is the effectiveness of the vaccine, we take $V = \zeta\tilde{V}$, and we adjust μ to capture this effectiveness factor.

Our epidemiological model is an augmented version of the standard *SIRS* model. We depart from the basic formulation in allowing for vaccination, V , in phase-II and the possibility of exogenous infection as represented by π .

For a large set of values (μ, ϕ) , the epidemiological model has a unique (and locally stable) steady state. However, the linear approximation to the dynamical system near the steady state has complex roots. Thus, for a fixed policy (including the no-intervention policy described below) the epidemic displays multiple “waves” of decreasing severity.

For $\gamma = \pi = 0$ (no loss of immunity and no exogenous infection) the model is a model of an *epidemic* in the sense that, in the long run, $I_t = 0$. For positive values of γ and/or π , this

and Qiu (2020)) that the effects of these NPIs are quite heterogeneous in terms of consequences of the epidemic. However, at the level of aggregation in this model they correspond to an average of feasible combinations. Future work will deal with heterogeneity in policies (different ϕ)

⁹Strictly speaking, the parameter μ measures the rate at which individuals leave the susceptible state weighted by the effectiveness of the vaccine.

is a model of an *endemic* disease as infections are recurrent and in the steady state a nonzero fraction of the population is in the infectious state.¹⁰ However, for small values of (γ, π) the long run fraction infected can be made arbitrarily small and the endemic case is very close to the epidemic case.

Finally, we label the path of the epidemic when the policy is *no intervention* (formally this corresponds to $\phi_t = 1$ and $\mu_t = 0$) the uncontrolled case, and we denote the resulting path by (\hat{S}, \hat{I}) .

The number of **deaths averted** up until time T , G_T , under a policy $\{\phi_s, \mu_s\}$ is

$$G_T = \int_0^T (\hat{D}_s - D_s) ds,$$

where $D_s = D_s^1 + D_s^+$ are total deaths under a given policy and $\hat{D}_s = \hat{D}_s^1 + \hat{D}_s^+$ are the equivalent deaths under the no intervention policy.

The measure of **relative deaths**, N_T is

$$N_T = \int_0^T \frac{\hat{D}_s}{D_s} ds,$$

and the **cumulative output cost** (relative to the full employment case) is

$$O_T = \left(\frac{1}{T}\right) \int_0^T \phi_s (1 - \zeta I_s) ds.$$

By comparing G_T and O_T we can estimate the output cost per death averted.¹¹

3 Analysis of the Model

Since the problem faced by the planner in Phases I and II is different, we start by studying the optimal policy in Phase II. We then discuss Phase I.

3.1 Phase II

In this Phase vaccination is available and the planner's objective function is

$$F(S, I, V) = \max_{\{\phi_t\}\{\mu_t\}} \int_0^\infty e^{-\rho t} [u(\phi_t w(1 - \zeta I_t) - c_V(\mu_t(1 - V_t))) - \Delta(D_t^1, D_t^+)] dt, \quad (8)$$

¹⁰At the time of this writing it is not clear what the duration of immunity associated with COVID-19 is.

¹¹It is clear that our measure of output cost ignores many other consequences of drastic reductions in economic activity. For example, all the managerial human capital that is lost (or reallocated to less-profitable activities) is not included in our measure. Moreover, to the extent that a given policy has implication for non-COVID deaths (e.g., deaths of despair, depression), our measure also ignores this. Thus, we believe that our formulation is somewhat conservative.

subject to equations (5), (6) and (7) and $S_0 = S$, $I_0 = I$ and $V_0 = 0$. The maximization is subject to $0 \leq \phi_t \leq 1$ and $0 \leq \mu_t \leq \bar{\mu}$, where $\bar{\mu}$ is a measure of the economy's maximal speed to vaccinate the population.

The optimal stay-at-home policy depends on the difference of the marginal shadow values of the stocks of infectious and susceptibles. Formally, interior case is when $\phi \in (0, 1)$, and the optimal ϕ solves (details in Appendix 1)

$$\frac{u'(\phi w(1 - \zeta I) - c_V(\mu(1 - V)))(1 - \zeta I)}{2\beta\phi(1 - \zeta)SI} = (F_S - F_I). \quad (9)$$

For a given state (S, I, V) the left-hand-side is decreasing in ϕ . F_I measures the contribution of an additional infected individual to the value of the problem, which is negative, and so is F_S for some combinations of the state (S, I, V) . However, in all cases, $F_S - F_I > 0$ since 100% of the susceptibles are in the labor force but only $1 - \zeta$ of the infected. Why is it that the optimal policy depends on the difference $F_S - F_I$? The reason is simple: A decrease in ϕ decreases the rate of infection (a positive), but it also slows down the decrease in the stock of susceptibles (a bad) and the optimal policy depends on the relative strength of the two effects. Thus, $-F_I - (-F_S)$ is the net gain from lowering ϕ . This net gain must be equal to the marginal welfare cost that is captured by the left-hand-side of equation (9). Decreases in ϕ are associated with larger gains of reducing the rate of infection relative to the rate at which the susceptible population shrinks. This gap is large at the beginning of an epidemic, as the stock of susceptibles is high and the level of infectiousness is low.

Why is the planner flattening the curve in this version of the model? There are two reasons. First, given our assumption about the impact of reaching hospital capacity, the term $-\Delta(D^1, D^+)$ is a concave function of I . This implies that smoothing the time path of deaths results in utility gains. Second, the existence of a technology with a capacity constraint induces another form of concavity that, in turn, makes smoothing the path of I_t optimal. As either $\eta \rightarrow 0$ and/or $\bar{\mu} \rightarrow 0$, Phase I is permanent, which implies that the only reason for the planner to decrease output is to distribute fatalities over time to avoid the higher fatality rate associated with exceeding the hospital capacity. Put differently, we expect that the lower the value of (η, μ) the higher the value of ϕ since the second “life saving technology” (vaccine) is less valuable.

Under some conditions the model has a steady state. For sufficiently small values of (γ, π) (the rate at which immunity is lost in the population and the rate at which exogenous infections occur) the steady state displays no output loss ($\phi^* = 1$) and no vaccination ($\mu^* = 0$). We summarize this result in the following proposition.

Proposition 1 (Phase II: Steady State) *Assume that the utility function is strictly increasing and strictly concave and that the marginal cost of vaccination is positive even at zero (that is, $c'_V(0) > 0$); then, for a small-enough value of (γ, π) , there exists a steady state characterized*

by $\phi^* = 1$ and $\mu^* = 0$, and the epidemiological variables are

$$S^* = \left(1 + \frac{\kappa + \gamma}{\beta(1 - \zeta)} \frac{\gamma + \pi}{\pi}\right) - \sqrt{\left(1 + \frac{\kappa + \gamma}{\beta(1 - \zeta)} \frac{\gamma + \pi}{\pi}\right)^2 - 4 \frac{\kappa + \gamma}{\beta(1 - \zeta)}},$$

and

$$I^* = \frac{\gamma}{\kappa + \gamma} (1 - S^*).$$

Proof. See Appendix 1 ■

If (γ, π) are sufficiently large, the steady state is interior.

3.2 Phase I

In this Phase there is no vaccine. We assume (as in Alvarez et al. (2020)) that vaccines become available at the Poisson rate η . The planner's problem is

$$W(S, I) = \max_{\{\phi_t\}} E \left[\int_0^{T_\eta} e^{-\rho t} [u(\phi_t w L_t) - \Delta(D_t^1, D_t^+)] dt + e^{-\rho T_\eta} F(S_{T_\eta}, I_{T_\eta}, V = 0) \right],$$

where the expectation is taken over the distribution of the stopping time T_η , which gives the first time that the Poisson process jumps. The expected time until a vaccine becomes available is $1/\eta$.

It is interesting to study what happens in Phase I as $t \rightarrow \infty$ along a realization in which the vaccine never arrives. The following proposition summarizes this case.

Proposition 2 (Phase I: Pseudo Steady State) *The Phase I model has a steady state that coincides with the steady state in Phase II.*

Proof. See Appendix 1 ■

This result says that after a long enough period of time the availability of a vaccine does not have a large impact on the optimal policy. Thus, the social value of a vaccine—measured as the impact on the continuation value—decreases to zero as $t \rightarrow \infty$. This, of course, ignores future epidemics that can be averted using vaccines since it only applies to situations characterized by small (γ, π) . In our quantitative section we show that the decrease in the value of the vaccine (although not to zero necessarily) is a robust feature.

This result has some implications for how to finance a vaccine. If the winner receives a patent, the economic value of that patent—again in the case of one epidemic—goes to zero as time goes by because the epidemic is being controlled. Specifically, the convergence result implies that for a large T_η , the change in optimal policy is small. One consequence of this is that firms that have to allocate resources to produce a patent see their potential payoff decreasing as time goes by, and intuitively, this results in lower investment.

This suggests that financing a vaccine with a prize with a fixed value can potentially be a more efficient mechanism relative to a patent—at least in terms of inducing resources to be allocated—to produce a vaccine in a shorter period.

4 Quantitative Results

In this section we present some results for a quantitative version of the model. We take our time unit to be a week. We report the complete list of parameters used in Appendix 3. Here, we describe the more significant assumptions underlying our baseline case:

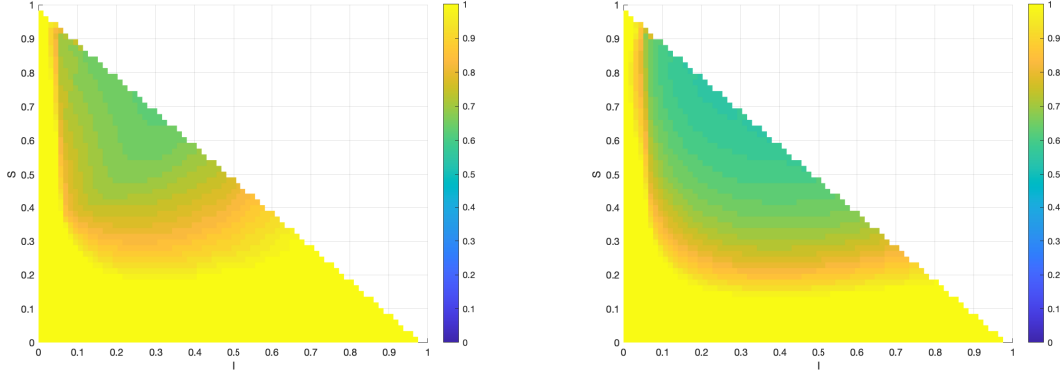
- \mathcal{R}_0 is 2.8, and for each diagnosed case there are approximately four asymptomatic cases ($\zeta = 0.25$).
- We assume that the case fatality rate (as a fraction of symptomatic cases) is slightly below 1.7%. (It is 1.8% in the U.S. at the time of this writing.) We assume that an individual remains infectious for about three weeks.
- We assume that 25% of the total hospital capacity can be allocated to COVID cases. This corresponds to 2.8×0.25 beds per thousand people.
- Vaccine: We assume that, in expectation, it takes about 50 weeks for a vaccine to become available (Phase II).¹² We also assume that the upper bound of the speed at which the population can be vaccinated is $\bar{\mu} = 0.036$. This is consistent with vaccinating 90% of the population in a year and a vaccine effectiveness of 0.85.

To compute the results we first discretized the continuous time HJB equation and then solved the weekly model using value function iteration. Given the model is highly non-linear, we solve the problem over a fine non-uniform grid and restrict the space to $0 \leq S + I + V \leq 1$. (See details in Appendix 2.)

4.1 Baseline Case

The optimal stay-at-home policies for each Phase are depicted in Figure 1. Panel (a) shows the optimal policy in phase I, ϕ^I , as a function of state space (S, I) . Panel (b) shows ϕ^{II} as a function of (S, I) for $V = 0$.

¹²This is $\eta = 1/50$.



(a) Phase I (ϕ^I)

(b) Phase II (ϕ^{II} at $V = 0$)

Figure 1. Optimal ϕ

Some interesting results include the following:

1. In Phase I (and Phase II) the optimal policy calls for no interventions for low levels of I . (The yellow area corresponds to $\phi = 1$.)
2. By construction, the optimal policy is a function of the state (S, I) (in Phase I). Any policy that chooses the severity of the stay-at-home policy considering only infectious (or deaths) is bound to be suboptimal.
3. The optimal policy in Phase II (when the vaccine is available and reported at $V = 0$, which is the state at the time of the switch) is slightly “shifted to the left” relative to Phase I, and it implies that, for all states, Phase II imposes more severe “stay-at-home” restrictions. In panel (b) the green-blue areas indicate more restricted choices. The reason for this is simple: If there is no vaccine ever, the optimal policy is close to “do nothing,”¹³ and it is only when there is something that can save lives (a vaccine in our setting) that it pays to restrict output to allow more individuals to get vaccinated.
4. There are large subsets of the state space that even if a vaccine is available it is optimal to restrict employment. For example, if 70% of the population is susceptible and 10% is infectious, then the optimal policy is to restrict employment to about 60% of the full employment level.

Any simulation must make an assumption about the realization of T_η , the time at which Phase II (vaccine) arrives. In our baseline we assume that $T_\eta = 50$, that is, that a vaccine becomes available after about 50 weeks, which is also the expected time of arrival.

Figure 2 shows the path of the relevant variables under the optimal policy. In the case of the stock of infectious individuals, we also include the value that it would take if the policy is $\phi_t = 1$ (referred as W/O intervention or the uncontrolled epidemic).

¹³Strictly speaking the model implies that $\phi < 1$ is optimal to distribute deaths over time in order to minimize the extra fatality rate associated with exceeding hospital capacity.

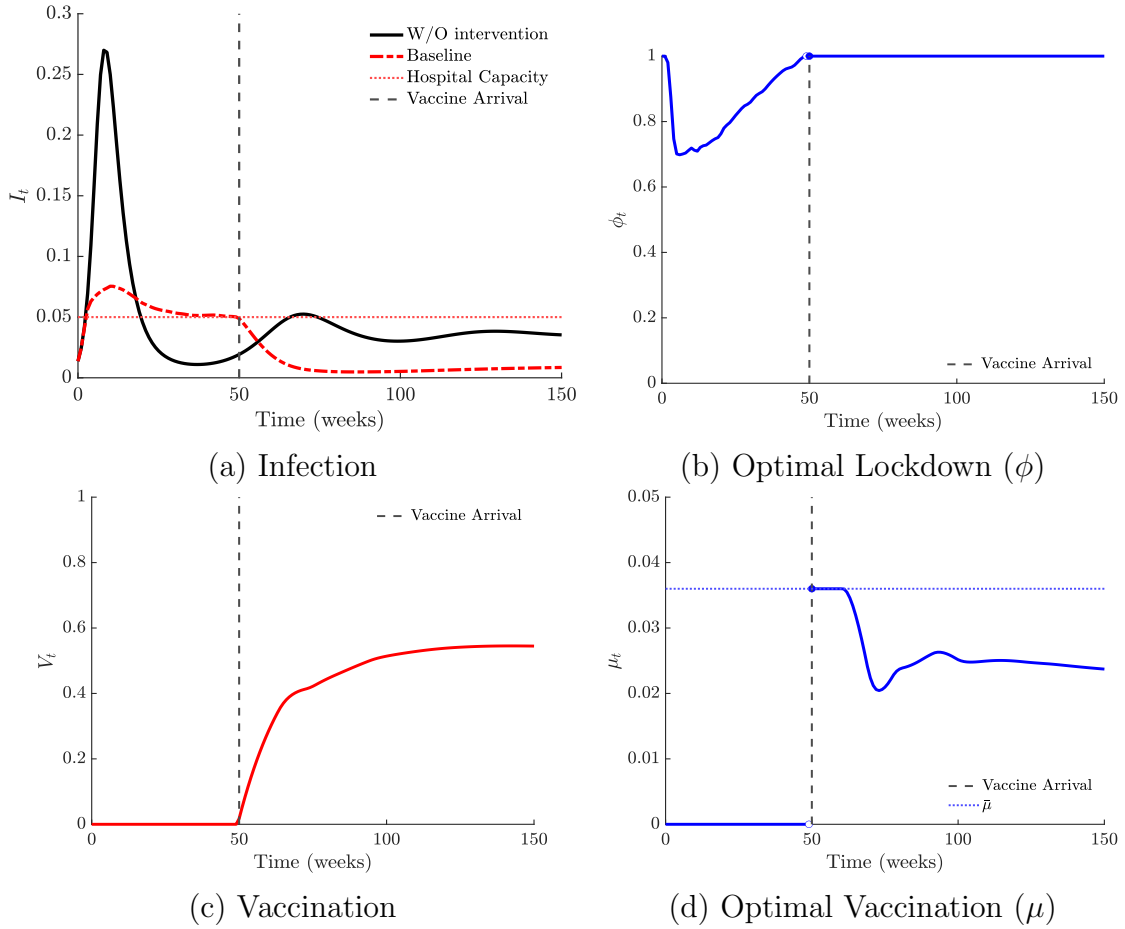


Figure 2: Baseline Paths

The results in the uncontrolled case are independent of the economic model and are driven by the assumptions embedded in the epidemiological model. In the absence of interventions/controls, the epidemic would peak at about 8 weeks, and about 27% of the population would be infectious at the time. The model implies that in the absence of any intervention there is a second wave that starts around week 40 and peaks in week 70 when 5.2% of the population is infected. In the absence of a policy (i.e., $\phi = 1$) there is a significant number of deaths.

Under the optimal policy the infectiousness curve (in red) is indeed flattened, and it takes about 10 weeks for the epidemic to peak. At the time of the peak, I is 7.5%. Under the optimal policy there is no second wave.

Panel (b) of Figure 2 displays the path of the optimal lockdown policy. Initially (in the first 2-3 weeks) there is a small lockdown that quickly increases, and it implies that employment is around 70% of the normal. Starting in week 12 there is a slow partial liberalization that ends in week 50 when all restrictions are lifted ($\phi = 1$). Thus, the optimal policy starts relaxing the stay-at-home constraint significantly **before** the peak of the epidemic. In this simple model the time path of ϕ is very close to the time path of output. Thus, output completely recovers in about a year.

Panel (c) shows the evolution of the stock of vaccinated individuals. Starting in week 50 the

optimal policy is to vaccinate as fast as possible (see Panel (d)), and the stock of vaccinated agents stays relatively constant at about 54%. The optimal vaccination rate has an almost bang-bang property to it: Initially it hits the upper bound and then rather sharply settles into its long-run value. In the long run it is optimal (for these parameter values) to continuously vaccinate individuals because of three factors: loss of immunity, vaccines that are not 100% effective and exogenous infections (this is quantitatively very small).

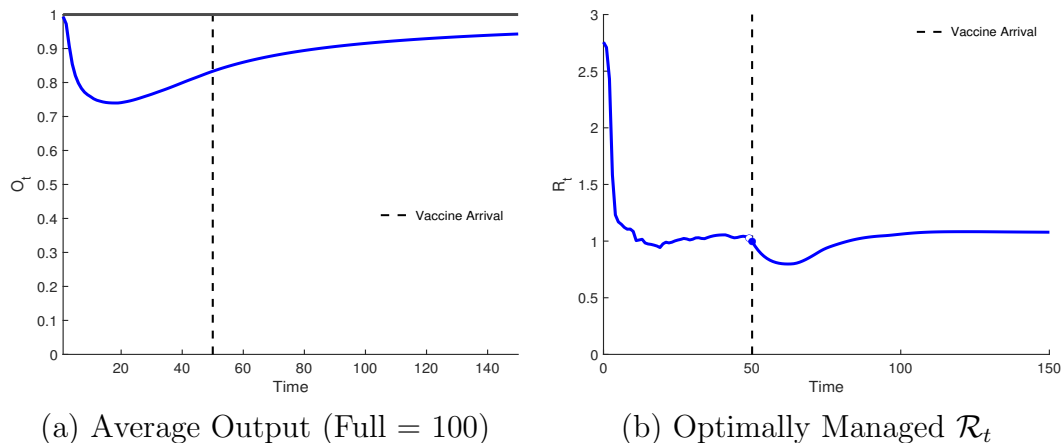


Figure 3. Baseline: Time Paths

Figure 3 shows the time paths corresponding to *average* output since the beginning of the epidemic (current output closely tracks ϕ_t) and the observed reproduction number. The largest drop in average output occurs in the first 20 weeks. After that point—which corresponds to higher values of ϕ —average output steadily increases. The output cost of the optimal policy is significant. For the first year the economy is operating on average at about 83% capacity. In the long run (about 7 years) the average output loss *over the whole period* is close to 2.5%. This is large.

Panel (b) of Figure 3 displays the time path of \mathcal{R}_t . In our calibration $\mathcal{R}_0 = 2.8$, but initially there is a large drop. It is 1.6 in week 3 and 1.10 in week 8. Starting in week 14 \mathcal{R}_t fluctuates around 1 with a slight upward trend that peaks at 1.05 in week 50. This result shows that simplistic policies (e.g., as advocated by Budish (2020)) that call for maximizing utility subject to the reproduction number being less than one are not necessarily optimal. The optimal policy is one where employment is increasing while the reproduction number is above one and getting higher.

The only random element in the model is the time at which a vaccine becomes available. It is not practical to report a large number of realizations, but it is interesting to discuss how the optimal policy should react if a vaccine is available earlier than expected. In the next section we describe one such realization.

4.1.1 Baseline: Early Arrival of a Vaccine

How does an early arrival of the vaccine change the optimal management of the epidemic? Figure 4 shows the outcome of our simulation for both the infection curve, I_t , and the optimal policy, ϕ , when the vaccine arrives fairly early, at about week 25 compared to the baseline case of week 50.

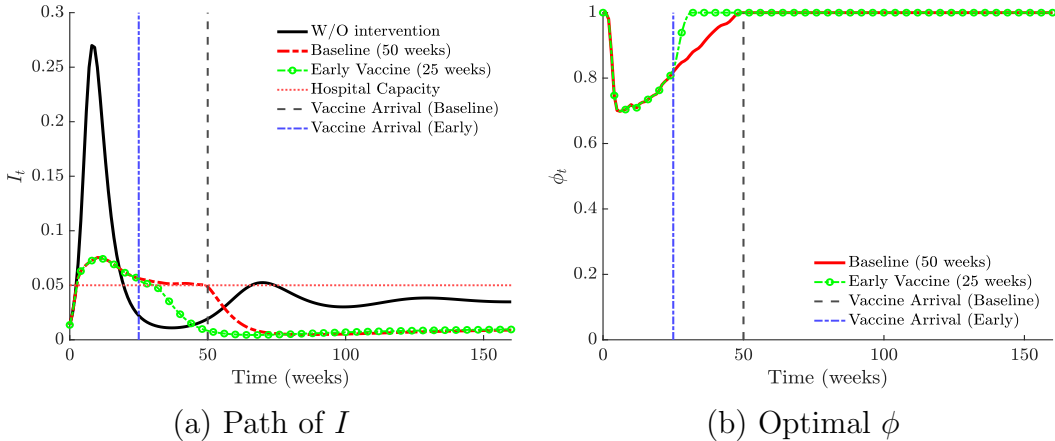


Figure 4: Early Arrival of a Vaccine ($T_\eta = 25$)

By construction the first 25 weeks display an optimal policy that is identical to the baseline. However, at that point the optimal policy (and the outcomes) differ significantly (the economy enters Phase II) relative to the baseline. The most salient changes include the following:

1. At the time the vaccine arrives the optimal stay-at-home policy starts a fast liberalization process, which is completed in about 5 weeks. The epidemic dies down a lot faster.
2. The number of averted deaths (not shown) is significantly higher than in the baseline and the cost lower: Output gets back to normal 7 weeks after the vaccine arrives.
3. The optimal vaccination policy is similar to the base case: Vaccinate at the highest feasible rate for about 5 months and then rather rapidly converge to the steady state.

In the model, availability of a vaccine plays a major role as it is the only technology that is available to save lives. Positive surprises (early arrival in our case) is equivalent to a large shock as measured by both the policy response and the consequences.

4.2 Uncertainty and Optimal Policies

How much faith should we put in our baseline results? One concern is that there is significant uncertainty both about how a lockdown will affect economic outcomes and about the appropriate values of many epidemiological parameters. This is an unfortunate situation but one that must be confronted head on. To convey a sense of how uncertainty affects the results, we

present the implications of the model when we vary several environmental variables. We find that the results are *extremely* sensitive to the underlying assumptions about the epidemiological model (κ and \mathcal{R}_0) and to some parameters that capture, in a rough sense, the quality of the health infrastructure (e.g., case fatality rate and vaccination speed).

Two Scenarios: Optimistic and Pessimistic We start the exploration of the role of uncertainty by considering deviations from the baseline in three dimensions: the case fatality rate, hospital capacity and the speed at which 95% of the population can be vaccinated.

We next describe the results for what we label optimistic and pessimistic scenarios.

- *Optimistic*: High vaccination rate (35% higher than in the baseline) and lower case fatality rate ($\chi = 0.25, \chi^H = 2.5\chi$) and 50% higher hospital capacity
- *Pessimistic*: Lower vaccination rate (35% lower than in the baseline) and higher case fatality rate ($\chi = 0.35, \chi^H = 2.5\chi$)

The results for the path of the epidemic and output are in Figure 5.

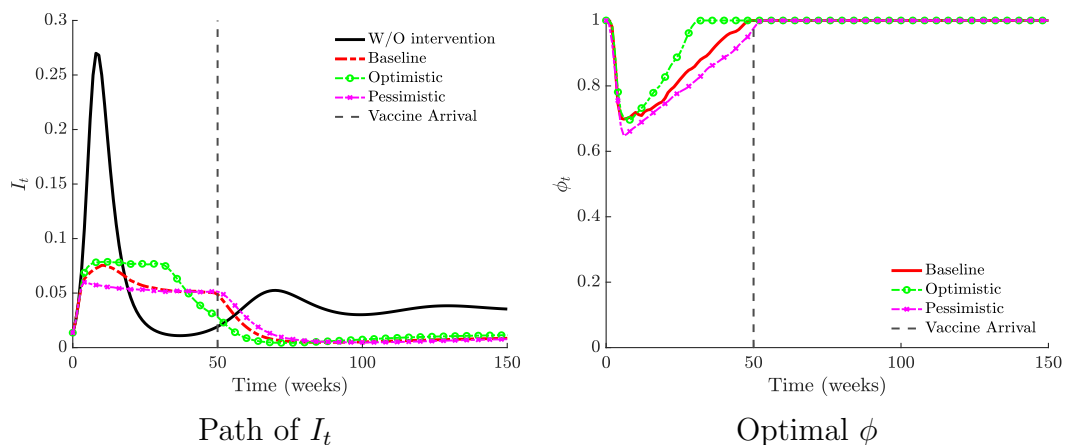


Figure 5. Optimistic and Pessimistic Scenarios

The differences are significant. At the peak of the epidemic the fraction of the population that is infectious can range from 7.8% (in the optimistic case) to 5.9% (in the pessimistic case). In the first 3-4 weeks the loss of output is somewhat similar in all three cases. However, in the optimistic scenario recovery occurs early and it is very fast. Monthly output is 100% of normal after 7 1/2 months. On the other hand, in the pessimistic case it takes over 12 months to reach that level.

Table 1 presents some summary statistics for the baseline and the two alternative scenarios. The differences in loss of output both in the short- and medium-run are large: In the pessimistic scenario the loss of output is twice as high as in the optimistic scenario. There is also a significant difference in the speed of recovery. In addition to the relevant economic variables, it shows the number of deaths averted and the output cost per death averted. The range is large both in terms of deaths averted (roughly 75% higher) and output cost per death averted (10% higher).

Table 1: Scenario Comparison			
Indicator	<i>Baseline</i>	<i>Optimistic</i>	<i>Pessimistic</i>
Y loss (1Y) (%)	17.4%	12.2%	19.9%
Y loss (3Y) (%)	6.9%	5.0%	10.8%
Full Recovery (months)	12	7.5	12
Deaths Averted (%) (2Y)	0.56%	0.40%	0.70%
Cost per Death Averted (\$) (2Y)	2.05M	2.09M	1.87M

The Impact of the Case Fatality Rate Our baseline assumes that the case fatality rate as a fraction of symptomatic cases is about 1.7%, and this corresponds to $\chi = 0.30$. We also explore the impact on the path of the epidemic and the optimal policy of two other values: $\chi = 0.25$ and $\chi = 0.35$, adjusting fatality rate when cases exceed hospital capacity as $\chi^H = 2.5\chi$. These correspond to case fatality rates of 1.4% and 2.5%, respectively. Figure 6 shows the results.

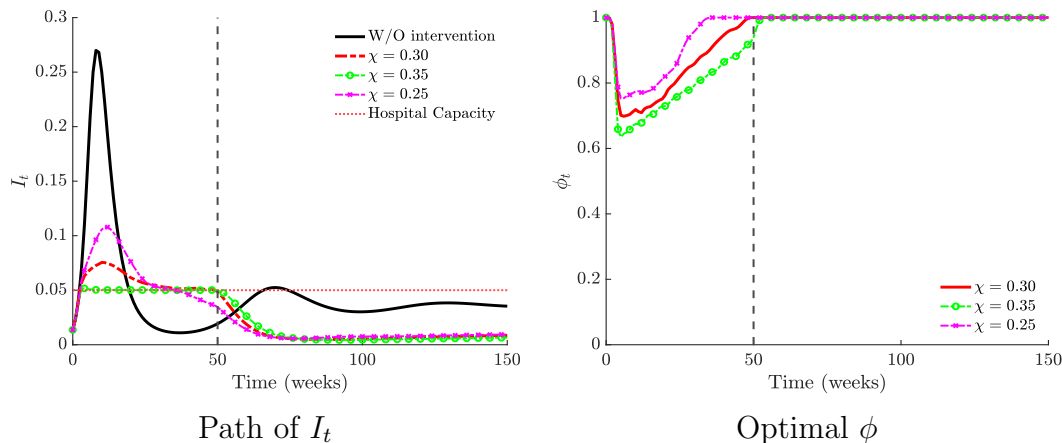


Figure 6. The Impact of the Fatality Rate

The lower the fatality rate the less aggressive the policy: The epidemic is short-lived and the economy recovers rapidly. The model implies that more-lethal epidemics result in more-severe lockdowns as the human costs increases.

The Impact of the Speed of Vaccination Figure 7 displays the time paths of infections, the optimal policy for the base case ($\bar{\mu} = 0.036$), and a very low vaccination capacity ($\bar{\mu} = 0.001$).

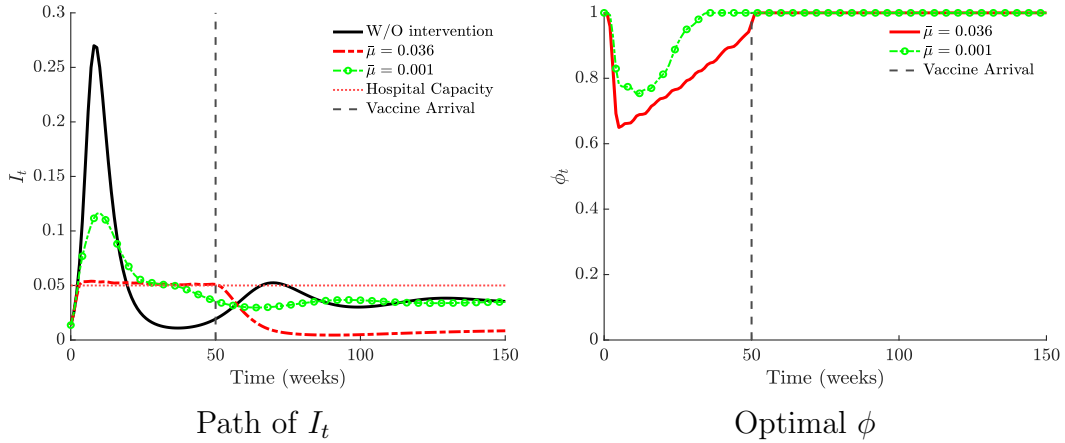


Figure 7. The Impact of Speed of Vaccination

The message is clear: When the ability to vaccinate is low (low $\bar{\mu}$) the optimal lockdown policy before the vaccine arrives imposes fewer restrictions. The reason, as discussed before, is that a low $\bar{\mu}$ implies the mitigation technology is not very productive and, hence, that the planner should not give up a lot of output today waiting for a technology that, even when it arrives, will not provide a fast solution. Thus, to the extent that $\bar{\mu}$ captures some dimension of the public health infrastructure implies that countries with poorer health infrastructure should adopt **less-restrictive** policies.

The Duration of the Infectious Period Figure 8 shows the effect of shortening the infectiousness period (two and a half vs. three weeks).

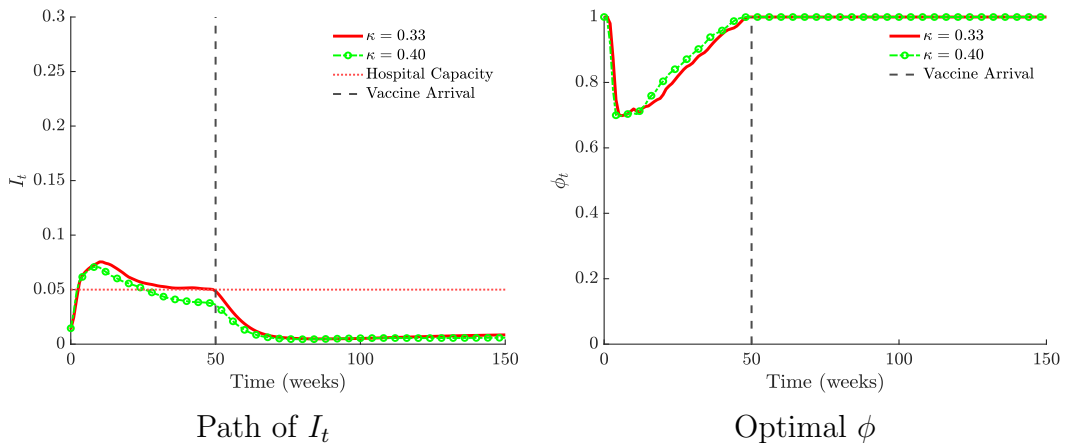


Figure 8. The Impact of the Duration of the Infectious Period

An epidemic characterized by a shorter infectious period has the property that, in the absence of controls, reaches a given fraction of the population faster. (In Figure 8 this corresponds to the green curve, $\kappa = 0.40$.) This, in turn, implies that the optimal response is to initially apply a slightly more severe lockdown policy and lift it sooner.¹⁴

¹⁴Even though the lockdown in the low κ case is more severe (lower ϕ) the speed at which the epidemic

The Reinfection Rate Our model assumes that individuals lose immunity at a rate given by γ . In the baseline the expected duration of immunity is about 50 weeks. In this section we report the effect on the optimal policy of a much longer immunity period (over 76 years) that approximates permanent immunity.

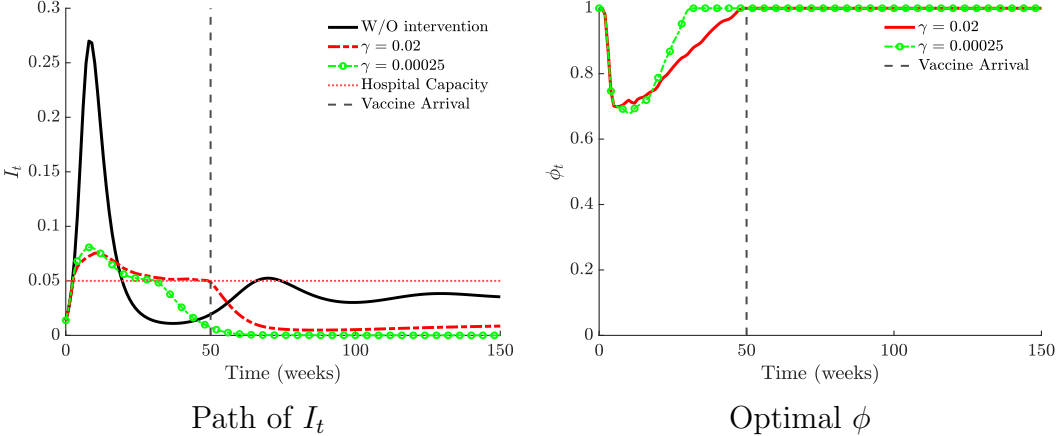


Figure 9. The Impact of Rate of Reinfection

Figure 9 shows the path of infection and the optimal lockdown policy in these two cases. Longer immunity does not have a large impact on the level of lockdown in the first few weeks, but it implies that the economy recovers sooner. It also implies that the level of infections is lower. This is driven by the transitory nature of the epidemic when the rate at which the population loses immunity is very small.

Hospital Capacity In the baseline we assume that 25% of the beds per thousand people (2.8×0.25) are available for COVID patients. On a per capita basis this implies that $\bar{H} = 0.05$. In this section we report the results associated with a 50% increase in hospital capacity.


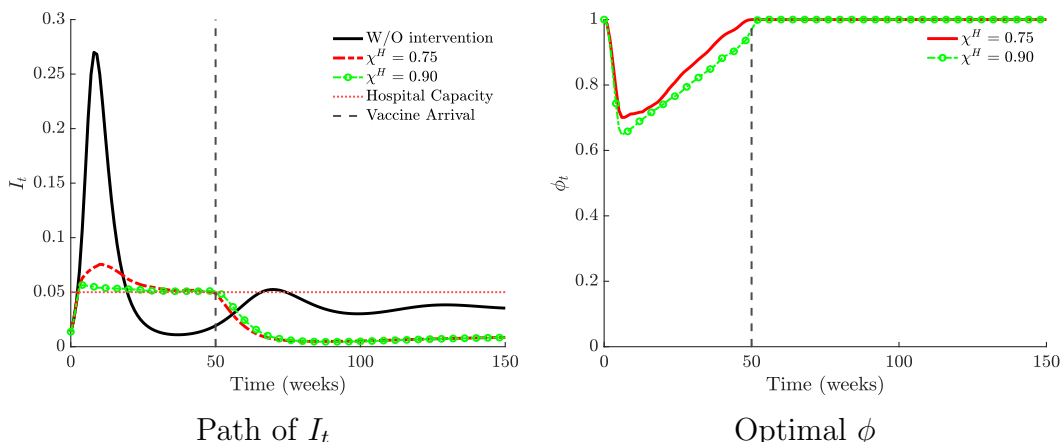


Figure 10. The Impact of Hospital Capacity

Higher hospital capacity results in smaller lockdowns (and of shorter duration). The level of infections is higher.

evolves is higher: The ratio of R_t paths ($R_t(0.4)/R_t(0.33)$) is greater than one. Thus, the planner does not fully compensate for the faster movement of the epidemic.

Case Fatality Rates When Hospital Capacity is Exceeded The variable χ^H measures the case fatality rate when the number of COVID cases exceed hospital capacity. The baseline χ is 0.30 and the baseline $\chi^H = 0.75$. We also consider an alternative $\chi^H = 0.9$. The results are in Figure 11.

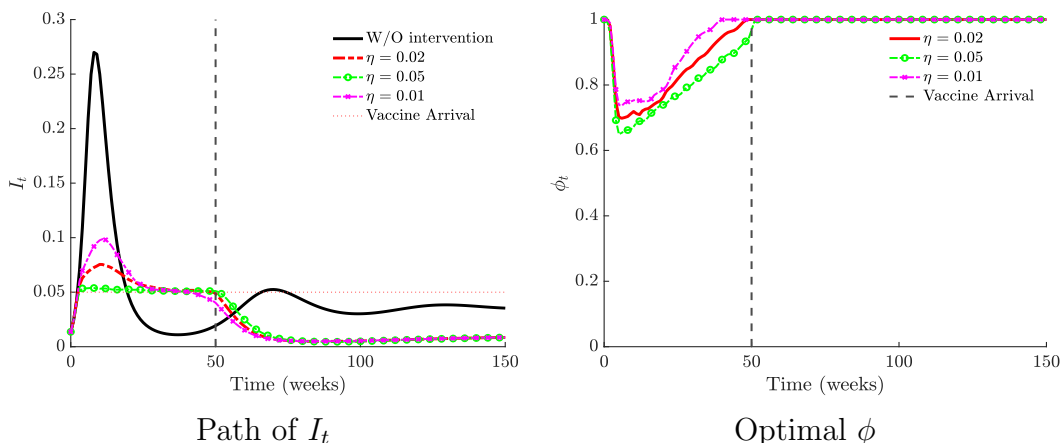


Path of I_t Optimal ϕ

Figure 11. The Impact of Case Fatality Rates

Higher fatality rates above hospital capacity result in a policy that keeps cases at the point where they do not exceed hospital capacity. This requires more-severe lockdowns and a lower infections during the first year.

Expected Time Until a Vaccine is Available In the baseline the expected time until a vaccine arrives is 50 weeks. We report the results in two alternative cases: fast vaccine arrival (20 weeks) and slow vaccine arrival (100 weeks).



Path of I_t Optimal ϕ

Figure 12. The Impact of Expected Time to Develop a Vaccine

When the vaccine is expected to arrive soon, the optimal policy is more stringent. The lockdown is more severe and infections and fatalities are lower. Why? The sooner the vaccine is available the sooner the planner will have a means of avoiding deaths. In this case, it pays to wait and keep fatalities low until the vaccine arrives. If the vaccine will take (in expectation) a long time to arrive, the cost in forgone output becomes too high and the optimal policy is to increase employment and output.

The Shadow Price of Life In the baseline we assume that the planner uses $v = 20$, which implies that each life is valued at about \$11 million. This figure is arbitrary but close to what is currently used by Federal agencies in project evaluation. Here we report the effect of changing the value of life to one quarter of a million ($v = 5$) and \$17 million ($v = 30$).




Figure 13. The Impact of the Shadow Price of Life

As expected, the higher the value that society puts on a human life the more severe the lockdown and the smaller the fraction of people infected.

Taking Stock Our sensitivity exercises are not extreme. The values of the key parameters are within the range of possible values about which, at this time, there is considerable uncertainty. Nevertheless, the policy implications for these cases differ significantly in terms of the key variable: how strict and how long is the optimal lockdown.

There is a clear message that emerges: policy makers face a difficult task in terms of choosing the optimal lockdown policy and knowledge of the relevant epidemiological and public health parameters is critical.

4.3 Optimal Vaccination Rate

In the previous comparative exercises we emphasized how the optimal lockdown policy responds to changes in certain parameters. In this section we report the impact of some changes in the basic environment on the optimal path of vaccination. We look only at changes that can potentially induce significant deviations from the base case. Figure 4 summarizes our findings.

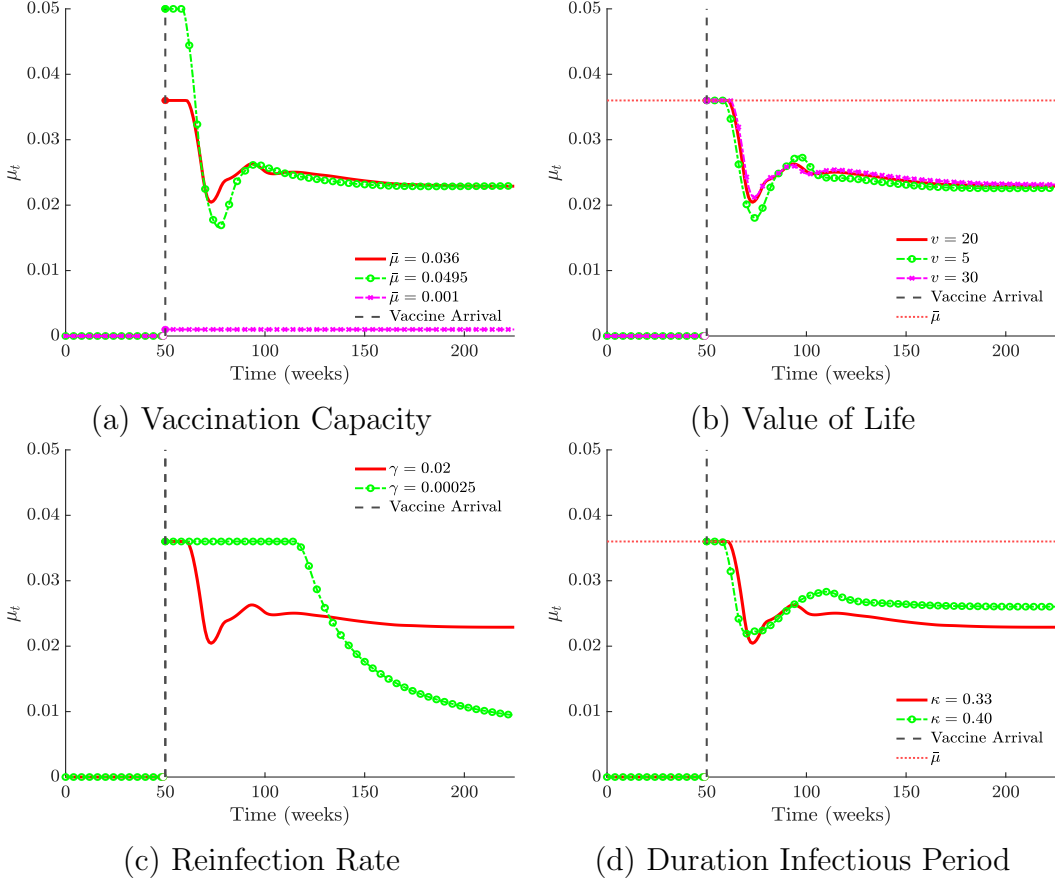


Figure 14: Optimal Vaccination Path

Overall, the optimal path of vaccination can be simply described as the following: Vaccinate as fast as possible and then settle to the long run (steady state) value. Differences in vaccination capacity have no impact on the long run (unless the capacity is binding). Panel (b) shows the impact of the value of life. The differences are small. Since the optimal policy initially is “vaccinate as fast as possible” the value of life has no impact. In the long run the vaccination rate is lower only when the value of life is very low (\$0.25M vs \$11M in the baseline).

Lower reinfection rates (Panel (c)) imply that vaccination is more effective as fewer people lose their immunity. This implies that the optimal policy makes a higher vaccination effort, and this implies that it hits the upper bound for a much longer period of time and then sharply decreases, to the steady state. Finally, Panel (d) shows the optimal vaccination when the infectious period is lower. As in the case of a lower γ , this increases the effectiveness of the vaccine and results in a much more aggressive vaccination policy.

4.4 The Value of Life and the Cost of Averting Death

In our base formulation of the social value of a life,

$$\Delta(D_t^1, D_t^+) = M_0 [k_a D_t^1 + k_e D_t^+]$$

where M_0 corresponds to the utility equivalent of the value of a statistical life.

In this section we present the results of several different values of a statistical life (M_0 in our notation). The value that we use in our baseline assumes that the present value of consumption of an individual who dies from the virus is \$11.25 million. We also consider several other cases indexed by the equivalent present value of consumption obtained from using a multiple of earnings, as a way of measuring other monetary costs of death (labeled PV in Table 2).

Table 2: The Impact of the Value of Life

PV (Mill)	Deaths Ave.	Cost (M) (2Y)	Y Loss (1Y) (%)	Trough (months)
\$0.26	0.52%	1.71	13.0	1
\$11.25	0.56%	2.05	17.4	1.5
\$16.99	0.60%	2.08	18.7	2

The results show that the valuation of life has a first-order effect on the optimal policy and, moreover, that the effects are highly nonlinear (using the present value of forgone consumption as the metric).

1. The higher the value of life the higher the loss of output and the slower the recovery.
2. The number of deaths averted (and the cost per death averted) vary as well. In particular, the cost per death averted in the high valuation case is about 22% higher than in the low case.
3. The optimal policy cannot be defined using the value of life. Put differently, even in the case when a life is valued at \$11.25 million it is optimal for the economy to spend only \$2 million to save a life.

4.5 The Social Value of A Vaccine

In order to evaluate the gains from a vaccine we compare the value (for each state (S, I)) of being in Phase II (with a vaccine available) relative to the case when no vaccine is available. Proposition 2 shows that even if a vaccine never arrives, the economy in a “permanent” Phase I converges to the same steady state as in Phase II when the parameters (γ, π) are small. In this case, if the difference between the total utility in each Phase is computed at a long enough horizon, our theory predicts that the value of a vaccine is small. The interesting quantitative question is to assess how much this social value changes for relatively short time horizons and for other parameters.

Table 3 reports the dollar equivalent of the gains associated with the availability of a vaccine for the baseline and the optimistic and pessimistic scenarios.

Table 3: Value of a Vaccine (Trillion)			
	Arrival Time (weeks)		
Scenarios	1	25	50
Baseline	3.16	2.06	1.34
Optimistic	2.10	1.07	0.87
Pessimistic	3.66	2.56	1.77

We include the gains at 3 points in time: at the beginning of the epidemic, 6 months into the epidemic and a year after the beginning of the epidemic. As expected there is a monotonicity corresponding to the results. The value of the vaccine increases as the environment gets worse. The absolute values are large and range from \$2.1 T to 3\$.66T. This is roughly 10% of the U.S. GDP (in the base case). In all three cases there is a significant decrease in the value of a vaccine —about 45% of the original value—that becomes available after one year. This decrease is driven by the change in the epidemic: More individuals are immune and hence the social value of a vaccine is lower.

Table 4 contains the same information except that it is indexed by the present value of consumption.

Table 4: Value of a Vaccine (Trillion)				
		Arrival Time		
v	PV (Mill)	1	25	50
5	\$0.26	2.63	1.29	1.00
20	\$11.25	3.16	2.06	1.34
30	\$16.99	3.30	2.30	1.41

The higher the value that society puts on a human life the more valuable the vaccine. In all three cases the value after a year is about 35% of the initial value.

What does this say about how to finance a vaccine? Our model is too stylized to discuss details of market failure. The values that we report correspond to a social valuation of a life. We find that to the extent the market values lives in a similar way, the market value of a patent decreases over time. (Table 4 gives options and reasonable bounds.) Our model predicts that if private firms engage in a discovery “race,” then, as the prize gets smaller, the less-promising candidate vaccines will be left out. Overall the private incentives decrease as time goes by if the source of revenue is a patent. To the extent that the social value differs from our estimates, and particularly if it does not decrease, then offering prizes instead of patents seems an efficient way of providing incentives to the private sector to produce a vaccine.

5 Concluding Comments

The goal of this paper is to analyze the optimal choice of stay-at-home policies in Phase I and stay-at-home and vaccination policy in Phase II. Optimal policies have a shock treatment aspect to them: strict lock-down initially with gradual liberalization that occurs before the peak of the epidemic. Similarly, the optimal vaccination policy is to vaccinate at capacity initially before converging to a steady state determined by the effectiveness of the vaccine or reinfection rate. Our novel finding is that pre-vaccination policies depend heavily on what happens after vaccine arrival (vaccination rate) and the probability of the arrival of a vaccine, and thus any analysis that does not model the post-vaccine phase in a realistic manner would recommend incorrect policies. The market value of a vaccine decreases rapidly—especially if the infection curve is not flattened under the optimal policy. The value that society places in individual life has a first-order effect on the optimal stay-at-home policy. This, in turn, determines the severity and duration of the recession associated with the epidemic. In many of our scenarios, the ex-post cost of averting a death is large and in the baseline case exceeds \$2 million. Our results illustrate that uncertainty about some features of the environment (e.g., the parameters that define the epidemiological model and the public health infrastructure) have a large impact on the optimal policy. Thus, producing better estimates of these key elements should be a priority in the policy area.

References

- [1] Adda, J. (2016), “Economic activity and the spread of viral diseases: Evidence from high frequency data,” *The Quarterly Journal of Economics*, 131(2), 891-941.
- [2] Acemoglu, D., V. Chernozhukov, I. Werning and M. Whinston, (2020), “A Multi-Risk SIR Model with Optimally Targeted Lockdown,” mimeo.
- [3] Alvarez, F., D. Argente and F. Lippi, (2020), “A Simple Planning Problem for COVID-19 Lockdown.” NBER working paper No 27102.
- [4] Atkeson, A. (2020), “What Will Be the Economic Impact of COVID-19 in the U.S.? Rough Estimates of Disease Scenarios,” Federal Reserve Bank of Minneapolis, Staff Report No. 595.
- [5] Aum, S., Y. Lee, and Y. Shin (2020), “Inequality of Fear and Self-Quarantine: Is There a Trade-off between GDP and Public Health?,” Working Paper.
- [6] Avery, C., W. Bossert, A. Clark, G. Ellison, and S.F. Ellison (2020), “Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists,” NBER Working Paper No 27007..
- [7] Baron, D, (2020), “Serology is Harder than It Looks,” Blog. April 14.
- [8] Baker, S. R., N. Bloom, S.J. Davis, K.J. Kost, M. C. Sammon, and T. Viratyosin (2020), “The Unprecedented Stock Market Impact of COVID-19,” NBER Working Paper No. 26945.
- [9] Barro, R.J., J.F. Ursúa, and J. Weng (2020), “The Coronavirus and the Great Influenza Pandemic: Lessons from the “Spanish Flu” for the Coronavirus’s Potential Effects on Mortality and Economic Activity” NBER Working Paper No. 26866.
- [10] Bassetto, M. (2020), “Latent Covid-19 Diffusion in Italy: An Estimate from Italian Provinces,” Mimeo.
- [11] Bairoliya, N. and A. Imrohoroglu (2020), “Macroeconomic Consequences of Stay-At- Home Policies During the COVID-19 Pandemic,” Mimeo.
- [12] Blavin, F. and D. Arnos (2020), “Hospital Readiness for COVID-19: Analysis of Bed Capacity and How It Varies Across the Country,” Working Paper Urban Institute.
- [13] Berger, D.W, K.F. Herkenhoff, and S. Mongey (2020), “An SEIR Infectious Disease Model with Testing and Conditional Quarantine,” NBER Working Paper No 26901.
- [14] Bick, A. and A. Blandin (2020), “Real Time Labor Market Estimates During the 2020 Coronavirus Outbreak,” Mimeo.

- [15] Bisin, A. and P. Gottardi(2021), “Efficient policy interventions in an epidemic,” *Journal of Public Economics* 200.
- [16] Budish, E. (2020), “Maximize Utility Subject to $\mathcal{R} \leq 1$: A Simple Price-Theory Approach to Covid-19 Lockdown and Reopening Policy,” BFI working paper 2020-31.
- [17] Caballero, R. and A. Simpssek (2020), “ A Model of Asset Price Spirals and Aggregate Demand Amplification of a Covid-19 Shock,” Working Paper MIT.
- [18] Chang, R. and A. Velasco (2020), “Economic Policy Incentives to Preserve Lives and Livelihoods,” Working Paper.
- [19] Chari, V.V. and C. Phelan, (2020), “Test the Healthy and Gather Data we Don’t Already Have,” *Minneapolis Star and Tribune* (April 3)
- [20] Chen, X. and Z. Qiu, (2020), “Scenario Analysis of Non-Pharmaceutical Interventions on Global COVID-19 Transmissions,” working paper.
- [21] Correia, S., S. Luck, and E. Verner (2020), “ Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu,” Working Paper.
- [22] Eichenbaum, M.S., S. Rebelo, and M. Trabandt (2020a), “The Macroeconomics of Epidemics,” NBER Working Paper No 26882.
- [23] Eichenbaum, M.S., S. Rebelo, and M. Trabandt (2020b), “The Macroeconomics of Testing During Epidemics,” NBER Working Paper No 27104.
- [24] Dingel, J.I., and B. Neiman (2020), “How Many Jobs Can Be Done at Home?,” Working Paper Becker Friedman Institute.
- [25] Droz, L.A. and M.M. Tavares (2020), “Responding to COVID-19: A Note,” Working Paper Federal Reserve Bank of Philadelphia WP 20-14.
- [26] Piguillem, F. and L. Shi, (2020), “Optimal Covid-19 Quarantine and Testing Policies,” working paper.
- [27] Farboodi, Maryam, Gregor Jarosch, and Robert Shimer, “Internal and External Effects of Social Distancing in a Pandemic,” NBER Working Paper No 27059.
- [28] Fang, H., L. Wang, and Y. Yang (2020), “Human Mobility Restrictions and the Spread of the Novel Coronavirus (2019-nCoV) in China,” NBER Working Paper No. 26906.
- [29] Faria-e-Castro, (2020), “Fiscal Policy during a Pandemic,” Working Paper Federal Reserve Bank of St. Louis 2020-006A,

- [30] Ferguson, N. et al. (2020), “Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand”
- [31] Ferguson, N.M., D. Laydon, G. Nedjati-Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunuba, G. Cuomo-Dannenburg, A. Dighe, I. Dorigatti, H. Fu, K. Gaythorpe, W. Green, A. Hamlet, W. Hinsley, L. C. Okell, S. van Elsland, H. Thompson, R. Verity, H. Volz, E. Wang, Y. Wang, P. GT Walker, C. Walters, P. Winskill, C. Whittaker, C. A. Donnelly, S. Riley, and A. C. Ghani, (2020), “Impact of Non-pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand,” Technical Report, Imperial College London.
- [32] Fernández-Villaverde, J. and C.I. Jones (2020), “Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities,” NBER Working Paper No. 27128.
- [33] Flaxman et al., (2020), “Estimating the Number of Infections and the Impact of Non-Pharmaceutical Interventions on COVID-19 in 11 European Countries.”
- [34] Fornaro, L. and M. Wolf (2020), “Covid-19 Coronavirus and Macroeconomic Policy: Some Analytical Notes,” Working Paper CREI.
- [35] Glover, A., J. Heathcote, D. Krueger, and J.V. Ríos-Rull (2020), “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic,” NBER Working Paper No 27046.
- [36] Gonzalez-Eiras and Niepelt (2020), “On the Optimal ‘Lockdown’ During an Epidemic” mimeo
- [37] Greenstone, M. and V. Nigam (2020), “Does Social Distancing Matter?,” Working Paper Becker Friedman Institute No 2020-26.
- [38] Gourinchas, P.O. (2020), “Flattening the Pandemic and Recession Curves,” Working Paper.
- [39] Gormsen, N.J. and R. S.J. Koijen (2020), “Coronavirus: Impact on Stock Prices and Growth Expectations,” Working Paper.
- [40] Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2020), “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?,” NBER
- [41] Hall, R.E., C.I. Jones, and P. J. Klenow (2020), “Trading-Off Consumption and COVID-19 Deaths,” Working Paper Stanford University.
- [42] Hortacsu, A., J. Liu and T. Schweg, (2020), “Identification of the Fraction of Unreported Infections in Epidemics with a Known Epircenter: An Application to COVID-19,” BFI working paper No. 2020-37.

- [43] Hsiang, S., D. Allen, S. Annan-Phan, K. Bell, I. Bolliger, T. Chong, H. Druckenmiller, A. Hultgren, L. Y. Huang, E. Krasovich, P. Lau, J. Lee, E. Rolf, J. Tseng, and T. Wu. (2020), “The Effect of Large-Scale Anti-Contagion Policies on the Coronavirus (COVID-19) Pandemic,” Working Paper.
- [44] Jones, C. J., T. Philippon, and V. Venkateswaran (2020), “Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home,” NBER Working Paper No 26984.
- [45] Kapicka, M. and P. Rupert (2020), “Labor Markets during Pandemics,” Working Paper Laboratory for Aggregate and Economics Fluctuations (LAEF).
- [46] Korolev, I. (2020), “Identification and Estimation of the SEIRD Epidemic Model for COVID-19,” working paper
- [47] Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020), “Scarring Body and Mind: The Long-Term Belief-Scarring Effects of COVID-19,” Working Paper.
- [48] Kubinec, R. (2020), “A Retrospective Bayesian Model for Measuring Covariate Effects on Observed COVID-19 Test and Case Counts,” Working Paper.
- [49] Kurman, A. E. Lale, and L. Ta (2020), “The Impact of COVID-19 on U.S. Employment and Hours: Real-Time Estimates with Homebase Data,” Mimeo.
- [50] Krüger, D. H. Uhlig, and T. Xie (2020), “Macroeconomic Dynamics and Reallocation in an Epidemic,” Working Paper Becker Friedman Institute No 2020-43.
- [51] Manski, C. and F. Molinari (2020), “Estimating the COVID-19 Infection Rate: Anatomy of an Inference Problem,” Working Paper.
- [52] Marchant, R., N. Samia, O. Rosen, M. Tanner, and S. Cripps, (2020), “Learning as we Go: An Examination of the Statistical Accuracy of COVID-19 Daily Death Count Predictions,” working paper.
- [53] Moghadas, S et al., (2020), “Projecting Hospital Utilization During the COVID-19 Outbreaks in the United States,” Proceedings of the National Academy of Sciences (downloaded 04/3/2020).
- [54] Pindyck R.S (2020), “COVID-19 and the Welfare Effects of Reducing Contagion,” NBER Working Paper No. 27121.
- [55] Shao, P. (2020), “Impact of city and residential unit lockdowns on prevention and control of COVID-19,” Working Paper.
- [56] Stock, J.H. (2020), “Data Gaps and the Policy Response to the Novel Coronavirus James H. Stock,” NBER Working Paper No. 26902.

- [57] Velde, F. R. (2020), “What Happened to the US Economy During the 1918 Influenza Pandemic? A View Through High-Frequency Data,” Working Paper Federal Reserve Bank of Chicago 2020-11.
- [58] Wang, L., Y. Zhou, J. He, B. Zhu, F. Wang, L. Tang, M. Eisenberg, and P.X.K. Song (2020), “ An epidemiological forecast model and software assessing interventions on COVID-19 epidemic in China,” Working Paper.

6 Appendix 1: Proofs

Discussion of Phase II Optimal Policy Let the value function of this problem satisfy the following HJB equation:

$$\begin{aligned} \rho F(S, I, V) = & \max_{(0 \leq \phi \leq 1, 0 \leq \mu \leq \bar{\mu})} \{u(\phi w(1 - \zeta I) - c_V(\mu(1 - V))) - \Delta [D^1, D^+]\} \\ & + F_S [-\beta\phi^2(1 - \zeta)SI - (\mu + \gamma + \pi)S + \gamma] + F_I [\beta\phi^2(1 - \zeta)SI - (\kappa + \mu + \gamma)I + \pi S]. \end{aligned}$$

The associated Hamiltonian is

$$\begin{aligned} H^{II} = & \max_{(0 \leq \phi \leq 1, 0 \leq \mu \leq \bar{\mu})} u(\phi w(1 - \zeta I) - c_V(\mu(1 - V))) - \Delta [D^1, D^+] \\ & + \lambda_S [-\beta\phi^2(1 - \zeta)SI - (\mu + \gamma + \pi)S + \gamma] + \lambda_I [\beta\phi^2(1 - \zeta)SI - (\kappa + \mu + \gamma)I + \pi S] \\ & + \lambda_V (\mu(1 - V) - \gamma V) + \hat{\gamma}^\phi(1 - \phi) + \hat{\gamma}_+^I(\bar{\mu} - \mu) + \hat{\gamma}_-^I\mu. \end{aligned}$$

The FOC are standard and given by the static conditions

$$\begin{aligned} u'()w(1 - \zeta I) &= 2\beta\phi(1 - \zeta)SI(\lambda_S - \lambda_I) + \hat{\gamma}^\phi, \\ \hat{\gamma}^\phi(1 - \phi) &= 0 \end{aligned} \tag{10}$$

and

$$-u'()c'_V(\mu(1 - V)) + \lambda_V\mu(1 - V) = \lambda_S S + \lambda_I I - \hat{\gamma}_-^I + \hat{\gamma}_+^I, \tag{11}$$

where we omit the arguments in the utility function to keep the expression simple. The constraints imply that $\hat{\gamma}_+^I(\bar{\mu} - \mu) = 0$, and $\hat{\gamma}_-^I\mu = 0$.

In the interior case, that is, when $\phi \in (0, 1)$, given that $\lambda_S = F_S$ and $\lambda_I = F_I$, equation (10) can be written as,

$$\frac{u'(\phi w(1 - \zeta I) - c_V(\mu(1 - V)))w(1 - \zeta I)}{2\beta\phi(1 - \zeta)SI} = (F_S - F_I),$$

which corresponds to equation (9) in the text.

Proof. In using the Hamiltonian H^{II} , care must be taken since the function $\Delta [D^1, D^+]$ has a kink at the hospital capacity level. Since we are trying to establish the existence of a steady state with arbitrarily small levels of infectious individuals, we will take the derivative of the function $\Delta [D^1, D^+]$ as if the level is below the hospital capacity. Thus, the relevant derivative with respect to I is

$$\Delta' = M_0 k_a \chi h \zeta \kappa$$

Consider first Phase II. The relevant co-state variables evolve according to the following differential equations (again omitting the arguments in the function $u'()$)

$$\dot{\lambda}_S = (\rho + \mu + \gamma + \pi)\lambda_S + (\lambda_S - \lambda_I)\beta\phi^2(1 - \zeta)I - \lambda_I\pi,$$

$$\begin{aligned}\dot{\lambda}_I &= u'()w\phi\zeta + \Delta' \\ &+ (\rho + \kappa + \mu + \gamma)\lambda_I + (\lambda_S - \lambda_I)\beta\phi^2(1 - \zeta)S, \\ \dot{\lambda}_V &= (\rho + \mu + \gamma)\lambda_V - u'()c'_V\mu\end{aligned}\tag{12}$$

The static first-order conditions are equations (10) and (11).

We conjecture that there is a steady state such that $\phi^* = 1$ and $\mu^* = 0$ when (γ, π) are arbitrarily small. Consider first the steady-state values of the state variables when $\pi = 0$ under the conjectured optimal policy. In this steady state the epidemiological variables satisfy

$$\beta(1 - \zeta)S^* = \kappa + \gamma,\tag{13}$$

and

$$I^* = \frac{\gamma}{\gamma + \kappa}(1 - S^*).\tag{14}$$

It suffices to show that the system of equations that is implied by $\dot{\lambda}_S = \dot{\lambda}_I = \dot{\lambda}_V = 0$ has a solution evaluated at the candidate steady state and satisfies equations (10) and (11).

Simple calculations show that $\dot{\lambda}_S = \dot{\lambda}_I = \dot{\lambda}_V = 0$ imply

$$\lambda_S^* = \frac{\beta(1 - \zeta)I^*(u'()\zeta w + \Delta')}{\Lambda}\tag{15}$$

and

$$\lambda_I^* = -\frac{(\gamma + \rho + \beta(1 - \zeta)I^*)(u'()\zeta w + \Delta')}{\Lambda},\tag{16}$$

$$\lambda_V^* = 0\tag{17}$$

where

$$\Lambda = \rho(\gamma + \rho + \beta(1 - \zeta)I^*) + \beta(1 - \zeta)I^*(\gamma + \kappa).$$

To complete the argument, it suffices to show that equations (10) and (11) hold as inequalities (ignoring the Lagrange multipliers). Some standard manipulations show that this is equivalent (in the case of equation (10)) to

$$\frac{u'()(1 - \zeta\frac{\gamma}{\gamma + \kappa}(1 - S^*))}{u'()w\phi\zeta + \Delta'} > \frac{2\gamma}{\Lambda} \left[\rho + \gamma + 2\beta(1 - \zeta)\frac{\gamma}{\gamma + \kappa}(1 - S^*) \right].$$

Since for arbitrarily small γ the left-hand-side remains bounded away from zero (it is actually increasing) and the right-hand-side converges to zero; this conditions is satisfied for sufficiently small (γ, π) .

To check that equation (11) holds as well, it suffices to show that

$$-u'()c'_V() < \frac{(u'()\zeta w + \Delta')}{\Lambda} \frac{\gamma}{\gamma + \kappa}(1 - S^*) \left[(\kappa + \gamma) - \left(\gamma + \rho + \beta(1 - \zeta)\frac{\gamma}{\gamma + \kappa}(1 - S^*) \right) \right]\tag{18}$$

and this holds for sufficiently small γ .

Next we show that the same steady state is a rest point of the dynamical system associated with the optimal solution in Phase I. The Hamiltonian in this case is

$$H^I = u(\phi w(1 - \zeta I) - \Delta [(D^1, D^+) + \eta F(S, I) + \psi_S [-\beta\phi^2(1 - \zeta)SI - (\gamma + \pi)S + \gamma] + \psi_I [\beta\phi^2(1 - \zeta)SI - (\kappa + \gamma)I + \pi S]),$$

where μ is exogenously set equal to zero and the relevant discount factor during Phase I is $\rho + \eta$.

We assume that the function $F(S, I, V)$ is differentiable for small values of I , and we look at the limiting behavior of the relevant dynamical system along a path in which the Poisson counter never goes off under the assumption that in the limit $\phi^* = 1$. The steady state (again set $\pi = 0$) is such that

$$\psi_S^* = \frac{\eta(\rho + \eta)F_S^* - (\beta(1 - \zeta)I^* + \pi)(- (u'())w\zeta + \Delta') + \eta F_I^*}{\tilde{\Lambda}},$$

$$\psi_I^* = \frac{(\rho + \eta + \gamma + \beta(1 - \zeta)I^*)(- (u'())w\zeta + \Delta') + \eta F_S^*(\kappa + \gamma)}{\tilde{\Lambda}},$$

where

$$\tilde{\Lambda} = (\rho + \eta + \gamma + \beta(1 - \zeta)I^*)(\rho + \eta) + \beta(1 - \zeta)I^*(\gamma + \kappa),$$

and

$$F_S^* = \lambda_S^* \text{ and } F_I^* = \lambda_I^*.$$

The sufficient condition to guarantee that $\phi^* = 1$ is

$$u'()w(1 - \zeta I^*) > 2(\psi_S^* - \psi_I^*)\beta(1 - \zeta)S^*I^*, \quad (19)$$

where, as in Phase II,

$$S^* = \frac{\gamma + \kappa}{\beta(1 - \zeta)} \text{ and } I^* = \frac{\gamma}{\gamma + \kappa}(1 - S^*).$$

The left-hand-side of equation (19) remains positive (and strictly bounded away from zero) as γ goes to zero. The right-hand-side converges to zero since I^* goes to zero as $\gamma \rightarrow 0$. ■

Appendix 2: Computational Appendix

$$D_t^1 = \chi \kappa \zeta h I_t$$

$$D_t^+ = (\chi^H - \chi)(\kappa \zeta h I_t - \bar{H})^+$$

$$\Delta(D^1, D^+) = M_0 [k_a D_t^1 + k_e D_t^+]$$

Writing the HJB equations for the two phases,

$$\begin{aligned} \rho F(S, I, V) = & \max_{(0 \leq \phi \leq 1, 0 \leq \mu \leq \bar{\mu})} \{u(\phi w(1 - \zeta I) - c_V(\mu(1 - V))) - \Delta[(D^1, D^+)] \\ & + F_S [-\beta \phi^2(1 - \zeta)SI - (\mu + \gamma + \pi)S + \gamma S] + F_I [\beta \phi^2(1 - \zeta)SI - (\kappa + \mu + \gamma)I + \pi S] \\ & + F_V [\mu(1 - V) - \gamma V]\}. \end{aligned}$$

$$\begin{aligned} \rho J(S, I) = & \max_{(0 \leq \phi \leq 1)} \{u(\phi w(1 - \zeta I)) - \Delta[(D^1, D^+)] \\ & + J_S [-\beta \phi^2(1 - \zeta)SI + \gamma(1 - S)] + J_I [\beta \phi^2(1 - \zeta)SI - (\kappa + \gamma)I]\}. \end{aligned}$$

Discrete Version of HJB

We discretize the above HJB equations and solve the weekly model using value function iteration. Given the model is highly non-linear, we solve the problem over a fine grid and restrict the interest region to $0 \leq S + I \leq 1$.

Phase II

$$\begin{aligned} F(S_t, I_t, V_t) = & \max_{(0 \leq \phi \leq 1, 0 \leq \mu \leq \bar{\mu})} \left\{ \frac{(1 - e^{-\rho \Delta})}{\rho} \left(u(\phi w(1 - \zeta I) - c_V(\mu(1 - V))) - \Delta(I_t) \right) + \right. \\ & \left. e^{-\rho \Delta} F(S_{t+\Delta}, I_{t+\Delta}, V_{t+\Delta}) \right\} \\ S_{t+\Delta} = & S_t + [-\beta \phi^2(1 - \zeta)SI - (\mu + \gamma + \pi)S + \gamma S] \Delta \\ I_{t+\Delta} = & I_t + [\beta \phi^2(1 - \zeta)SI - (\kappa + \mu + \gamma)I + \pi S] \Delta \\ V_{t+\Delta} = & V_t + [\mu(1 - V) - \gamma V] \Delta \end{aligned}$$

Phase I

$$\begin{aligned} J(S_t, I_t) = & \max_{(0 \leq \phi \leq 1)} \left\{ \frac{(1 - (e^{-(\rho+\eta)\Delta})}{\rho + \eta} \left(u(\phi w(1 - \zeta I)) - \Delta(I_t) + \right. \right. \\ & \left. \left. \eta F(S_t, I_t, 0) \right) + e^{-(\rho+\eta)\Delta} J(S_{t+\Delta}, I_{t+\Delta}) \right\} \\ S_{t+\Delta} = & S_t + [-\beta \phi^2(1 - \zeta)SI - (\gamma + \pi)S + \gamma S] \Delta \\ I_{t+\Delta} = & I_t + [\beta \phi^2(1 - \zeta)SI - (\kappa + \gamma)I + \pi S] \Delta \end{aligned}$$

Appendix 3: Calibration

1. **Utility** We consider log utility. We also assume that minimum consumption, $\underline{c} = (1-c_0)w$. Hence,

$$u(\phi w(1 - \zeta I) - c_V(\mu(1 - V))) = \ln[\phi wL - c_V(\mu(1 - V)) - (1 - c_0)w].$$

If there is no vaccination and $\phi = L = 1$, $1 - c_0$ is the fraction of steady-state output that captures the minimum level of consumption.

We assume that $c_0 = 0.6$ and, hence, that output cannot fall below 40% of its steady-state value, if the vaccination is low.

2. **Vaccination**

- (a) *The cost of vaccination:*

$$c_V(\mu(S + (1 - \zeta)I)) = c_V^0(\mu(S + (1 - \zeta)I))^{1+c_V^1}.$$

Baseline: $c_V^0 = 0.1w$, $c_V^1 = 0$.

- (b) *Vaccination capacity* We set $\bar{\mu} = 0.036$

3. **Epidemiological Parameters**

- (a) *The Fraction ζ .* A difficulty estimating ζ is the lack of random testing at this point and the ad-hoc assumptions about mortality that have to be made to produce estimates. Hortacsu et al. (2020) estimate a range for ζ . Their results—based on data prior to the institution of stay-at-home policies in many states in early March 2020—imply that $\zeta \in [0.4, 0.25]$. Li et al. (2020) use a different approach and rely on Chinese data estimate $\zeta = 0.04$. More recent evidence give a range between 0.10 and 0.30. We assume $\zeta = 0.25$.
- (b) *The Fatality Rate.* What are the implied case fatality rates? We assume that the flow of deaths is given by $D_t = \chi\kappa h\zeta I_t$. There are several assumptions underlying this definition. First, we assume that the individuals who die from the flow out of the infectious category (κI_t) must have shown symptoms (hence the flow is $\kappa\zeta I_t$) and were hospitalized (the flow is $\kappa h\zeta I_t$). For this subset the fatality rate is χ . Then the fatality rate for the overall symptomatic population is χh . In the US at the time of this writing, the measured death rate is about 2%, while for the world as a whole it exceeds 3%. At the same time there are some countries—including several Latin American countries—in which the case fatality rate is below 1%. Thus, it seems that $\chi = 0.3$ given a hospitalization rate of $h = 0.056$ is a reasonable estimate. This implies a measured case-fatality rate of 1.68%.

- (c) *The Recovery Rate κ* . If, on average, individuals exit the infected category (to either resistant or deceased) in 3 weeks, then $\kappa = 1/3 = 0.33$.
- (d) *The Gross Transmission Rate β* . We view estimates of \mathcal{R}_0 as more reliable than estimates of β . Our strategy is to use estimates of \mathcal{R}_0 to estimate β . In our base case $\mathcal{R}_0 = 2.8$. Then given

$$\frac{\beta(1 - \zeta)}{\kappa} = \mathcal{R}_0,$$

we estimate β as

$$\beta = \frac{\kappa \times \mathcal{R}_0}{1 - \zeta} = \frac{0.33 \times 2.8}{0.75} = 1.232.$$

There is significant uncertainty about the relevant value of \mathcal{R}_0 . Many studies put the range of \mathcal{R}_0 between 1.5 and 4.0. A recent study by Fernandez-Villaverde and Jones (2020) that matches the evidence with the SIR model—but that imposes an arbitrary sequence ϕ_t —estimates that $\mathcal{R}_0 = 4.2$ and even higher in some European countries.

Baseline: $\beta = 1.1$

4. Economic/Institutional Parameters

- (a) *Output per worker* Our unit of analysis is an individual. We assume that there are 328×10^6 individuals and GDP of \$20 trillion/year. Thus, output per worker per week is \$1,173:

$$w = 1,173.$$

- (b) *Discount factor* We assume that the annual discount factor ρ is somewhere between 1 and 3%. The base case is $\rho = 0.0123$ on an annual basis. Since the model is weekly we have that

$$\rho = 0.000236.$$

This value has the “property” that the present discounted value of weekly output of the average worker (who earns twice as much as the average person since only 50% of the population work) satisfies

$$\frac{2,331}{0.000233} = 10,000,000$$

which is not an unreasonable number.

5. The Δ function.

- (a) *Base Case: Estimation of M_0* . The constant is given by

$$\ln(vw - \underline{c}) \frac{1 - e^{-\rho T}}{\rho}.$$

In the baseline, we assume that $v = 20$ and the implied present value of consumption (using the same life expectancy) is \$11.25M, which is on the high side. Thus we have the following values:

Present Value of Consumption	v	M_0
16,990,933	30	5116
11,250,753	20	4914
2,640,482	5	4205

6. **Definition:** \mathcal{R}_t

$$\mathcal{R}_t = \frac{\beta(1 - \zeta)\phi_t^2 S_t}{\kappa}.$$

7. Next we pick γ . In our baseline, $\gamma = 0.02$, i.e., 2% of the recovered + vaccinated become susceptible in unit time. This is likely an upper bound based on the evidence so far, so we also consider $\gamma = 0.00025$ in the exercise.

Table 1: Baseline Parameters

Meaning	Parameter	Value
Fraction Diagnosed among Infected	ζ	0.25
Immunity Loss Rate	γ	0.02
Basic Reproduction Number	R_0	2.8
Recovery Rate	κ	0.33
Discount Rate	ρ	0.000233
Time Step	Δ	1/5
Loss function	v	20
Output per Worker	w	1173
Mortality parameter	χ	0.3
Mortality due to Overwhelming of Hospitals	χ^H	2.5 χ
Fraction Hospitalized	h	0.056
Hospital Duration	κ^H	1/3
Hospital Capacity	\bar{H}	0.05
Vaccine Cost	c_v	0.1 w
Vaccination capacity	μ	0.036
Minimum Consumption (1- c_0)	c_0	0.60
Vaccine Arrival Poisson	η	1/50
Exogenous Source of Infection	π	0.0005
Initial I_0	I_0	0.01
Initial S_0	S_0	0.99

Appendix 4: Behavioral Model

In this section, we outline a simple model that captures the behavioral response to individual risks associated with the infection (see, for e.g. Bisin and Gottardi (2021) for a model with externalities). Individuals differ along three dimensions: i) their risk assessment while working, $\sigma(i)$, ii) their risk assessment while not working, $\nu(i)$, with $\nu(i) < \sigma(i)$ as working increases the probability of infection, and iii) their disutility cost from losing a job, $\lambda(i)$. Workers choose whether or not to work, and we allow for some workers to be forced to not work to capture the impact of closures in activities where there is no option of working from home (e.g., restaurants).

There are three types of workers: those who a) work, b) choose not to work, and c) are forced to not work, with measure π_t . We assume that everyone has the same consumption and same private perception of risk, $H(D, D^+)$, which depends on the aggregate flow of deaths and deaths due to hospital capacity. Flow utility of three types of workers is as follows:

$$\begin{aligned} \text{a) work:} & \quad u(c - \bar{c}) - \sigma(i)H(D^1, D^+) \\ \text{b) choose not to work:} & \quad u(c - \bar{c}) - \lambda(i) - \nu(i)H(D^1, D^+) \\ \text{c) forced to not work:} & \quad u(c - \bar{c}) - \nu(i)H(D^1, D^+) \end{aligned}$$

Let the measure of agents be denoted by K . The measure of individuals who can work is $(1 - \pi_t)$. Among the individuals who can work, individuals who choose to work satisfy

$$\begin{aligned} \frac{\lambda(i)}{\sigma(i) - \nu(i)} & \geq H(D^1, D^+) \\ \text{Let } A(D^1, D^+) & \equiv \left\{ i \text{ s.t. } \frac{\lambda(i)}{\sigma(i) - \nu(i)} \geq H(D^1, D^+) \right\} \end{aligned}$$

Let $A^c(D^1, D^+)$ be the complement of $A(D^1, D^+)$.

Thus,

$$\phi_t = (1 - \pi_t)K(A(D^1, D^+))$$

If we assume that the government implements a transfer policy that equalizes consumption across the three different types, aggregate preferences are given by

$$\begin{aligned} & u(c_t - \bar{c}) - \left[(1 - \pi_t) \int_{A(D_t^1, D_t^+)} \sigma(i) dK(i) \right. \\ & \left. + (1 - \pi_t) \int_{A^c(D^1, D^+)} \nu(i) dK(i) + \pi_t \int_{A(D_t^1, D_t^+) \cup A^c(D^1, D^+)} \nu(i) dK(i) \right] H(D_t^1, D_t^+) \\ & \quad - \int_{A^c(D^1, D^+)} \lambda(i) dK(i) \end{aligned}$$

Therefore, the loss function can be interpreted as

$$\Delta(D_t^1, D_t^+) = \left[(1 - \pi_t) \int_{A(D_t^1, D_t^+)} \sigma(i) dK(i) + (1 - \pi_t) \int_{A^c} \nu(i) dK(i) + \pi_t \int_{A(D_t^1, D_t^+) \cup A^c} \nu(i) dK(i) \right] \\ H(D_t^1, D_t^+) + \int_{A^c} \lambda(i) dK(i),$$

and this utilitarian objective function coincides with the social preferences that we use.

Appendix 5: Literature Review

We are fully aware that this summary is incomplete and that closely related research is probably missing. However, we still want to acknowledge some of the recent work in the topic, and found useful to group different papers in different categories, with the obvious caveat that these are imperfect and that there exists a substantial amount of overlap in the existing work.

Optimal management of the epidemic in SIR models: These papers use optimal control techniques to explore the management of an epidemic. See Alvarez, Argente, and Lippi (2020), Gonzalez-Eiras and Niepelt (2020), Acemoglu, Chernozhukov, Werning, and Whinston (2020), Jones, Philippon, and Venkateswaran (2020), Piguillem and Shi (2020).

Policies in an SIR models: These papers explore the effects of different policies in the dynamics of the epidemic (i.e., social distancing, lockdown, etc...) in SIR models developed by Kermack and McKendrick (1927). See for example, Atkeson (2020), Berger, Herkenhoff, and Mongey (2020), Neumeyer (2020), Bassetto (2020), Droz and Tavares (2020), Hsiang et al. (2020), Fang, Wang, and Yang (2020), Shao (2020), Wang et al. (2020). Avery, Bossert, Clark, Ellison, and Ellison (2020), Farboodi, Jarosch, and Shimer (2020), Aum, Lee, and Shin (2020). Glover, Heathcote, Krueger, and Ríos-Rull (2020), Pindyck (2020), Chang and Velasco (2020).

Measurement issues: These papers discuss measurement issues in the data or key parameters on SIR model (Stock 2020, Korolev, 2020, Kubinec, 2020, Manski and Molinari 2020, Fernández-Villaverde and Jones 2020, Hortacsu, Liu and Schwieg 2020, Harris 2020, Blavin and Arnos, 2020, Greenstone and Nigam 2020, Hall, Jones, and Klenow 2020).

Macroeconomic implications of epidemics: Some of the research provides a historical perspective by analyzing the economic implications of past pandemics as Adda (2016), Barro, Ursúa, and Weng (2020), Correia, Luck, and Verner (2020), Velde (2020). A number of papers explores macro implications of epidemic shocks and policy interventions. See Eichenbaum, Rebelo, and Trabandt (2020a,b), Fornaro and Wolf (2020), Bairoliya and Imrohorglu (2020), Krüger, Uhlig, and Xie (2020), Kozlowski, Veldkamp, Venkateswaran (2020), Bodenstein, Corsetti, and Guerreri (2020), Guerrieri, Lorenzoni, Straub, and Werning (2020), Faria-e-Castro (2020), Caballero and Simpssek (2020), as well as implications for different markets (i.e., labor market Bick and Blandin 2020, Kapicka and Rupert 2020, Dingel and Neiman 2020, Kurman, Lale, and Ta 2020, stock market i.e., Alfaro, Chari, Greenland, and Schott 2020, Baker, Bloom, Davis, Kost, Sammon, and Viratyosin 2020, Gormsen, N. and Koijen 2020).