

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

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

First Version: April 27, 2020

This Version: June 3, 2020
[\[Most Updated Version Here\]](#)

Abstract

We study a dynamic macro model to capture the trade-off between policies that simultaneously decrease output and the rate of transmission of an epidemic. We find that optimal policies initially restrict employment but partial loosening occurs **before** the peak of the epidemic. The arrival of a vaccine (even if only a small fraction can be vaccinated in the short run) implies a relaxation of stay-at-home policies and, in some cases, results in an increase in the speed of infection. The monetary value of producing a vaccine decreases rapidly as time passes. The value that society assigns to averting deaths is a major determinant of the optimal policy.

*The views expressed in this paper do not necessarily reflect the position of the Federal Reserve Bank of St. Louis. We thank Emilio Espino, Matt Famiglietti, Miguel Faria-e-Castro, Fernando Navajas, Andy Neumeyer 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 and NBER

[§]Washington University in St. Louis

1 Introduction

The objective of this research is to understand how features of the economy and the parameters describing an epidemic influence the choice of policy during an epidemic. We consider a government that can impose restrictions on employment along the lines of “stay-at-home” policies, and can allocate resources to attain a certain level of vaccination when a vaccine becomes available.

The macro literature on the impact of pandemics is large and growing. This paper is closest to the recent work of Alvarez, Argente, and Lippi (2020), Acemoglu et. al (2020) and Gonzalez-Eiras and Neipelt (2020). The major differences with Alvarez et. al. are 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, and that we are able 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.¹

The details of how a vaccine interacts with other policies is novel and interesting. We find that unless 100% of the population can be instantaneously vaccinated—a patently unrealistic case—arrival of a vaccine 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, 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.

We posit that individual preferences depend on individual consumption—ignoring the private value of life—and that social preferences take into account the utility loss associated with deaths. We discuss the consequences of using different values of a statistical life since this turns out to be important in our quantitative exercises. On the epidemiological dimension, we use a standard SIR model with (endogenously chosen) vaccination rates in some states and less than perfect immunity.

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 we 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.

¹Acemoglu et. al. study optimal lockdown for heterogenous agents and find that the optimal policy call for different lockdown strategies for different individuals. Gonzalez-Eiras and Neipelt present a general model but they concentrate of special cases in order to find closed form solutions. Appendix 5 has a partial list of the rapidly growing literature on the economic effects of COVID-19.

For a developed country like the U.S. we find that it is in general optimal to vaccinate at the highest possible rate (institutionally determined). For that reason, in the quantitative section we simply set the vaccination cost to zero (which implies that the optimal policy is to vaccinate at the highest feasible rate). In the case of less developed countries the cost of vaccination is not trivial and, in those cases, our theoretical model provides guidance.

On the theoretical side, we show that the model has a steady state and, more interesting, that along a path in which a vaccine or a treatment never becomes available² (Phase I) 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 (ignoring recurrences) 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 calibrate the model using standard estimates of the epidemiological parameters and we find that optimal policies are very sensitive to the details of the model, about which there is significant uncertainty. Some of our more interesting findings include:

1. The optimal policy depends on both the number of infected **and** susceptible individuals.³ In most of our simulations the optimal stay-at-home policy in Phase I (no vaccination or treatment available) implies:
 - (a) A **sharp decrease** in employment that ranges, depending on the particular case, between 20 and 35 percent. This reflects the high payoff of lowering the reproduction number when it is very high (i.e. when the fraction of the population susceptible is high) which mechanically occurs at the beginning of the epidemic.
 - (b) A gradual liberalization (e.g. allowing more economic activity to take place) that occurs **always** before the epidemic reaches a peak.
 - (c) A fairly **wide range** for the estimated lockdown time (i.e., a few weeks to several quarters)..
 - (d) If a vaccine arrives early (economy enters Phase II), the optimal response is a significant reduction of the restrictions on employment even if only a small proportion of the population can be vaccinated in a week. This liberalization of the lockdown is sometimes accompanied by an **increase** in the rate at which the virus spreads.⁴
2. Concavity of preferences implies that individuals prefer relatively constant consumption over time. A managed epidemic that flattens the infection curve lowers consumption at some points in time but it avoids the more pronounced peaks that would occur if no policy

²Although optimal policies take into account that the probability is positive.

³Implementation of the optimal policy requires random testing. The case for random testing has been made by many. Among economists a good discussion is in Chari and Phelan (2020).

⁴Whether this happens or not depends on whether the peak of the epidemic has been reached.

is put in place. However, in our quantitative exercises we find that this preference for smoothness has a small quantitative impact on the optimal solution. What has a large effect is the additional value of society puts on averting deaths. The assumed value of a statistical life has a first order impact on the economic performance.

3. We find that varying the value of a statistical life but keeping everything else constant (at the baseline parameters):
 - (a) The number of deaths averted as a fraction of the population can range from a low of 0.07% to a high of 0.39%.
 - (b) The cost per death averted if a vaccine becomes available in about a year, is in the range of 2.6 to 10 million. Moreover the policies that put the lowest value on human life have the costliest results in terms of output per life averted (49 million).
4. The social market value of a vaccine depends on the specific scenario (e.g. different epidemiological parameters and value of a statistical life). In general, we find that the rate of decrease in the value (measured in dollars) is large, and the value at the one year mark is a small fraction of the initial value. This result is robust to the possibility of a second wave.

We illustrate evaluate the model in a large number of scenarios, and we show that the results are *extremely* sensitive to the underlying assumptions about the epidemiological model as well as some parameters that capture, in a rough sense, the quality of the health infrastructure (i.e., case fatality rate and vaccination speed) and the value of life. It seems that a good way to convey our concern about how uncertainty, but at the same time it provides bounds on the optimal management of the epidemic. Even when we restrict the analysis to reasonable bounds we find widely varying implications.

In section 2 we present the model and in section 3 we discuss some theoretical results. In section 4 we present our quantitative findings. Section 5 briefly discusses ongoing work on extensions and section 6 offers 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” restriction 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 enter 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 preferences 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 on the availability of a vaccine:

- *Phase I*: This is the period in which there is no vaccine available. The only available tool is “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 like 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 an exogenous maximum) the speed at which the population is vaccinated.⁵

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 $u(\phi wL - c_V(\mu(S + (1 - \zeta)I)))$, where the second term captures the cost in terms of output of vaccinating a population of size $S + (1 - \zeta)I$. This is the population that includes susceptible and infected individuals who are asymptomatic. Of course, this term is operative only in Phase II when a vaccine is available, together with a bound on the speed at which the population can be vaccinated.

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 static social payoff is

$$u(\phi wL - c_V(\mu(S + (1 - \zeta)I))) - \Delta(D).$$

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 $\Delta(D)$ is increasing and convex.

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).

⁵Alvarez et. al. (2020) assume that most of the population can be instantaneously vaccinated.

Formally, preferences are given by

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

where the expectation is taken over the realization of T_η .

2.1.1 Special Case

The special case assumes no cost for the vaccine ($c_V = 0$), and the instantaneous payoff is

$$u(\phi w L) - \Delta(D) = N \ln(w\phi L - \underline{c}) - N(M_0 D).$$

In this formulation \underline{c} is the minimal level of consumption, N is population size, and 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 is

$$M_0 = \ln(vw - \underline{c}) \frac{1 - e^{-\rho T}}{\rho}.^6$$

In ongoing work we consider several variations that accommodate the idea that society is willing to spend more resources to avert “excess deaths” that is, deaths over and above some baseline.⁷

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.⁸ Here, we present a simple version although more general formulations (e.g. hospitalizations as a separate state with its own law of motion,

⁶There 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).

⁷It 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, creating “vaccination stations” in convenient places (e.g. supermarkets, 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.

⁸One 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)

alternative matching function to replace the canonical βSI in the SIR model) are relatively easy to incorporate.

In the model 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 $(1 - \zeta)I$.

Let S be the number of susceptible individuals and R the population of resistant individuals. Then the potential labor force, L , is given by

$$L = S + R + (1 - \zeta)I. \quad (2)$$

Since we normalized the population to one this is

$$L = 1 - \zeta I. \quad (3)$$

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.⁹

Finally we assume that a certain fraction of the resistant/recovered lose their immunity (at rate γ). The simple model is then given by

$$\begin{aligned} \dot{S} &= -\beta(\phi S)(\phi(1 - \zeta)I) - \mu S + \gamma(1 - S - I) \\ &= -\beta\phi^2(1 - \zeta)SI - \mu S + \gamma(1 - S - I). \end{aligned} \quad (4)$$

The first term is the standard matching function of the SIR model, while the second term, μS is the population that becomes resistant as a result of vaccination. The last term, $\gamma(1 - S - I)$ captures both the rate at which resistant individuals lose their immunity and the entrance of new susceptible individuals in the population.

The stock of infectious evolves according to

$$\dot{I} = \beta\phi^2(1 - \zeta)SI - \kappa I. \quad (5)$$

At this aggregate level this simple model suffices. However, keeping track of hospitalizations, deaths and individuals who have immunity (recovered if it turns out that infection provides immunity) is possible and desirable using a more disaggregated model.

In the simple model, we do not keep separate track of deaths associated with the epidemic. However, they play an important role determining the optimal policy. To keep the model simple

⁹In 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 and Qiu (2020)) that the effects of these NPIs is 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).

¹⁰It is interesting to note that in this simple version of the SIR model knowledge of I_t (the level of infections) and the rate of change over time (dI_t/dt) suffices to inform a planner that knows ϕ what S_t is.

—and ignoring the obvious lags— we will assume that a fixed fraction of those individuals who exit the infectious/symptomatic state, $\kappa\zeta I$, die. We denote this fraction by χ . Then, the flow of deaths at time t is $D_t = \chi\kappa\zeta I_t$.

Finally if we denote the path of the epidemic in the absence of a policy —what we label the uncontrolled case— by (\hat{S}, \hat{I}) . The number of **deaths averted** up until time T , G_T , under a policy $\{\phi_s, \mu_s\}$ is

$$G_T = \chi\kappa\zeta \int_0^T (\hat{I}_s - I_s) ds,$$

the measure of **relative deaths**, N_T as

$$N_T = \int_0^T \frac{\hat{I}_s}{I_s} ds,$$

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 discussing the optimal policy contingent on the economy having switched to Phase II first. We then discuss Phase I.

3.1 Phase II

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

$$F(S, I) = \max_{f\phi_t g f\mu_t g} \int_0^T e^{-\rho t} [u(\phi_t w(1 - \zeta I_t) - c_V(\mu_t (S_t + (1 - \zeta)I_t))) - \Delta(D_t)] dt, \quad (6)$$

subject to equations (4) and (5) and $S_0 = S$ and $I_0 = I$ and 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 speed to vaccinate the population.¹²

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

$$\frac{u'(\phi w(1 - \zeta I) - c_V(\mu(S + (1 - \zeta)I)))(1 - \zeta I)}{2\beta\phi(1 - \zeta)SI} = (F_S - F_I). \quad (7)$$

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.

¹²Although we assume that the planner knows the aggregate fractions of S and I , we do not assume knowledge of which individual is infected. Thus, the planner has to vaccinate all the asymptomatic individuals even though the vaccine is “wasted” on those infected.

For a given state (S, I) the left hand side is decreasing in ϕ . F_I measures the contribution of an additional infected individual to the value of the problem and it is negative, and so is F_S . 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 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 (7). 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? Given the linearity of the $\Delta(D)$ function and a small discount factor (in our calibration is about 1% per year) distributing deaths over time provides very little gain. The main motivation—given the absence of treatment—is to “buy time” until a vaccine becomes available.¹³

The model has a steady state. For sufficiently small γ (the rate at which immunity is lost in the population) the steady state displays no output loss ($\phi = 1$) and no vaccination ($\mu = 0$). We formally 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, $c_V^l(0) > 0$) then, for a small enough γ , there exists a steady state characterized by $\phi = 1$ and $\mu = 0$ and the epidemiological variables are*

$$S = \frac{\kappa}{\beta(1 - \zeta)},$$

and

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

Proof. See Appendix 1 ■

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

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

¹³In the extensions we discuss an alternative $\Delta(D)$ function that allows for differential death rates depending on the the number of cases relative to measures of hospital capacity. This gives two additional reasons for the planner to slow down the epidemic: avoiding higher death rates and since this view implies that the $\Delta(D)$ is concave there is a smoothing argument.

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 is discovered is $1/\eta$.

The key differences between phases I and II are the non-availability of a vaccine (in Phase I) and a higher discount factor (also in Phase I). It is interesting to study what happens in Phase I as $t \rightarrow \infty$ and there is no switch. 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.

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 T_η large 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 should result in fewer resources allocated to discovering a vaccine as time goes by.

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 the baseline. 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 that for each diagnosed case there are approximately six asymptomatic case ($\zeta = 0.15$)
- We assume that the infectious period lasts 3 weeks.
- Vaccine: We assume that, in expectation, it takes about 50 weeks for a vaccine to become available (Phase II)¹⁴. We also assume that it is costless to administer a vaccine and that the upper bound of the speed at which the population can be vaccinated is $\bar{\mu} = 0.10$.


¹⁴This is $\eta = 1/50$.

This implies that it takes 30 weeks to vaccinate 95% of the population from the time that a vaccine becomes available. In our numerical results we simplify and assume zero cost of administering the vaccine. This implies that $\mu_t = \bar{\mu}$ in Phase II which differs from the optimal policy (that has $\lim_{t \rightarrow \infty} \mu_t = 0$).¹⁵ Finally, we report the results optimistic and pessimistic bounds and we provide some bounds associated with uncertainty over the values of some parameters.

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 $0 \leq S + I \leq 1$. (See details in Appendix 2)

4.1 Baseline Case

The optimal policy ϕ as a function of population that is susceptible and infected (S, I) is displayed (for Phases I and II) in Figure 1. In either phase, the yellow area corresponds to no intervention ($\phi = 1$), whereas the orange to green colors represent different degree of intervention. Notice that for very low levels of I , the optimal policy calls for no intervention.



(a) Phase I

(b) Phase II

Figure 1. Optimal ϕ

The optimal policy illustrates some interesting results:

1. By construction the optimal policy is a function of the state (S, I) and it is such that for regions of the state space the optimal policy is not to restrict output (bright yellow area). Any policy that chooses the severity of the stay-at-home policy considering only infectious (or deaths) is bound to be suboptimal.¹⁶

¹⁵We solved a number of cases with endogenous μ but for low values of reinfection, γ , the solution is always bang-bang and, hence, setting it at the upper bound at zero cost is a very good approximation.

¹⁶As indicated before, using information on I_t and dI/dt it is possible to infer S_t .

2. The optimal policy in Phase II (when the vaccine is available) is slightly “shifted to the right” relative to Phase I and it implies that, for all states, Phase II imposes less severe “stay-at-home” restrictions.
3. There are large subsets of the state space that even if a vaccine is available it is optimal to restrict employment. In the model arrival of a vaccine is not equivalent to lifting restrictions. It depends on the state of the economy.

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 that path of the stock of infectious individuals, I , in two cases: uncontrolled epidemic (black) and optimally managed epidemic (red).

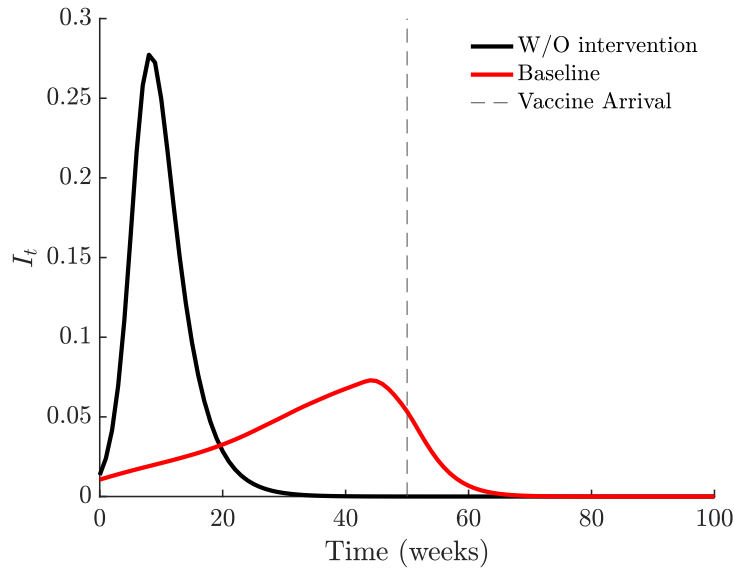


Figure 2. Time Path of the Epidemic

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 controls, the epidemic would peak at about 9 weeks and about 28% of the population would be infectious at the time. 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 a little less than a year for the epidemic to peak (44 weeks). At the time of the peak I is 7.3%.

Panel (a) of Figure 3 displays the path of the optimal policy in this realization while panel

(b) shows the implied \mathcal{R}_t .

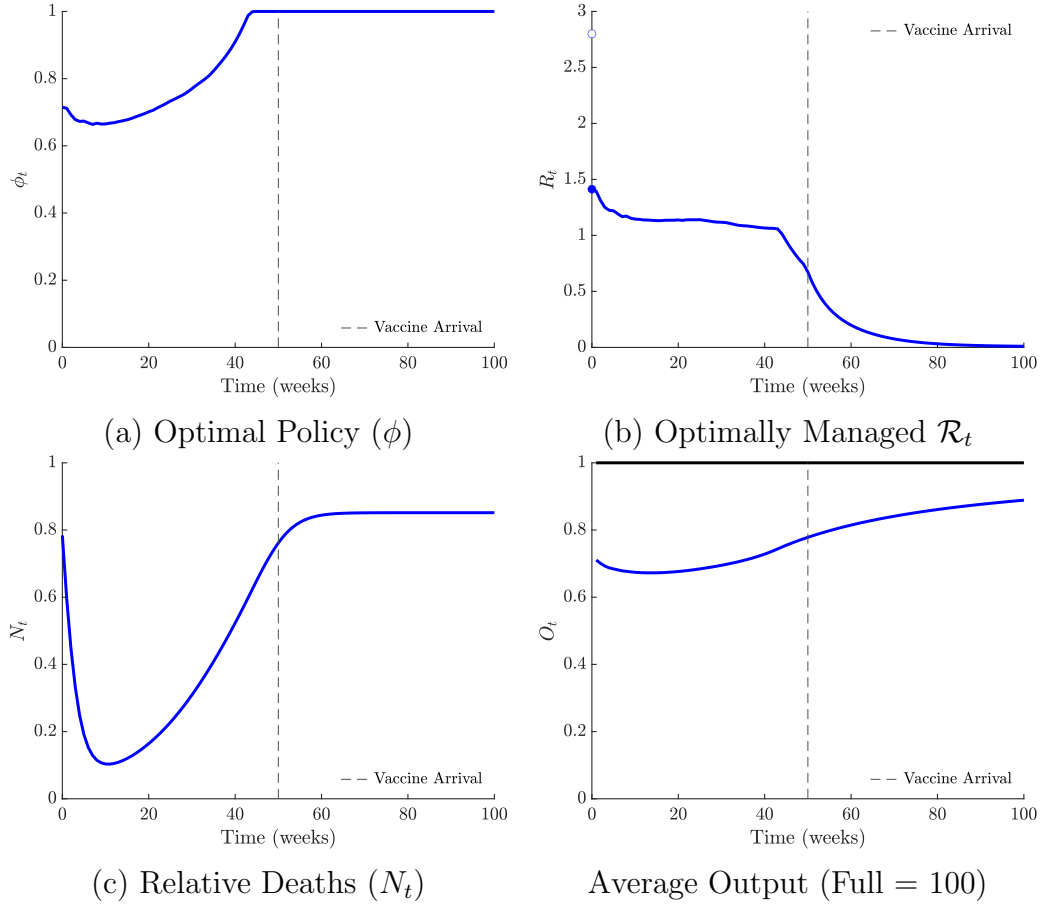


Figure 3. Baseline: Time Paths

Two features of the solution are worth emphasizing. First, the initial “stay-at-home” policy is fairly aggressive and employment (and output in our linear model) falls approximately 30-34% from its full employment level. Starting in week 8 there is a slow partial liberalization that ends in week 44 when all restrictions are lifted. Thus, the optimal policy relaxes the stay-at-home constraint significantly **before** the peak of the epidemic. Moreover, at the time of the peak all restrictions are eliminated even though there are new cases. In this simple model the time path of ϕ is very close to the time path of output. Thus, output completely recovers in about 11 months.¹⁷

In our calibration $\mathcal{R}_0 = 2.8$ but the first observed \mathcal{R}_t is 1.41 due to the aggressive restrictions on output that are implemented in the first week, and it shows large decreases for the first 6 weeks. After that time the optimal policy keeps \mathcal{R}_t slightly above one until the epidemic peaks.¹⁸

Panel (c) shows the implications of the model for relative deaths, N_t , and panel (d) for the *average* loss of output, O_t (relative to potential which is set equal to 100). Since the optimal

¹⁷In the Appendix we present time series of monthly output for a variety of cases.

¹⁸We present the time path for \mathcal{R}_t for several cases in Appendix 4.

policy flattens the the infection curve, relative deaths trends down at the beginning of the epidemic (few people infected and few die). In the long run, the optimal policy manages to keeps deaths at about 85% of the uncontrolled case. As mentioned before, in this version of the model, the only “technology” available to reduce the fatality rate is vaccination. Thus, flattening the curve is basically “waiting for a vaccine” that will decrease mortality.

The output cost of the optimal policy is significant. For the first year the economy is operating at about 78% capacity. This estimate is substantially smaller than what is found in other studies. In the long-run (about 150 weeks) the average output loss *over the whole period* exceeds 7%. This is large. It is straightforward to calculate the output cost per death averted. In this case, the economy gives up about 12.6 million in consumption to avoid one death.¹⁹

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

In Figure 4 we display 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.

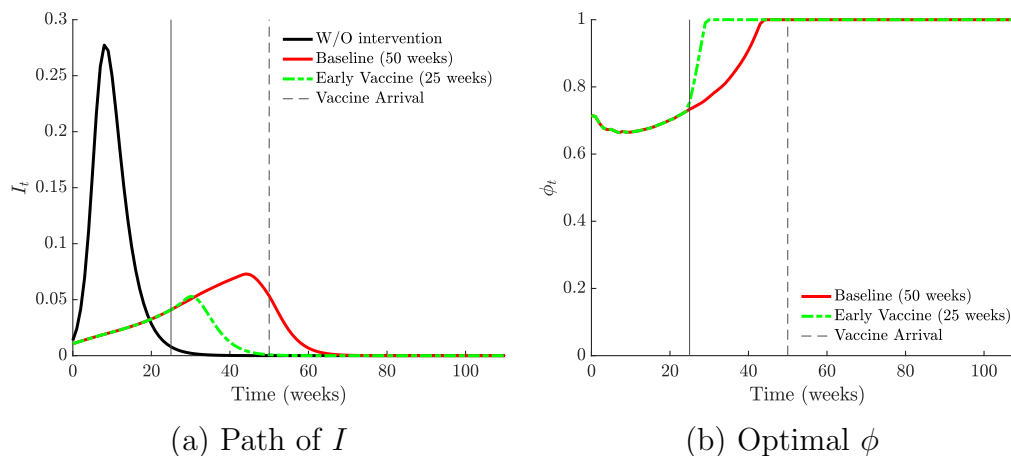


Figure 4: Early Arrival of a Vaccine

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 are:

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.

¹⁹In Appendix 4 we present the time path of the output cost per death averted for a variety of scenarios

2. This optimal policy implies a significant liberalization **before** the epidemic reaches a peak of infectious individuals (week 30).
3. A somewhat surprising consequence is that, following the arrival of the vaccine, the number of infectious individuals actually **grows faster** than before. The reason for this is that the availability of a vaccine increases the downward drift in the number of susceptible individuals (they are getting vaccinated) and, hence, it lowers the rate of contagion going forward. For the policy makers, the cost of the epidemic in terms of future deaths and consumption decreases (less future contagion) and, consequently, the marginal cost in terms of current output should decrease as well. The implied liberalization, but not full opening of the economy, implies more contact among individuals that, in turn, pushes up the contagion rate.
4. The number of averted deaths (not shown) is significantly higher than in the baseline and the cost lower: the ex-post output cost of averting one death is 2.5 million.

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 path of infections.

4.2 The Consequences of Uncertainty

In this section we provide some results corresponding to changes in parameters that capture, in a rough sense, the quality of the health infrastructure —as captured by the case fatality rate χ and the upper bound of the vaccination rate $\bar{\mu}$ — as well as the speed at which the epidemic unfolds as captured by the infectiousness parameter κ .

It seems that a good way to convey our concern about how uncertainty about epidemiological as well as public health parameters impact the results is to consider reasonable deviations of the baseline in two dimensions: the case fatality rate and speed at which 95% of the population can be vaccinated. In what follows we describe the results for what we label optimistic and pessimistic scenarios.

- *Optimistic*: High vaccination rate (95% of the population can be vaccinated in 12 weeks) and lower case fatality rate ($\chi = 0.04$).
- *Pessimistic*: Lower vaccination rate (95% of the population can be vaccinated in 60 weeks), and higher case fatality rate ($\chi = 0.06$).

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

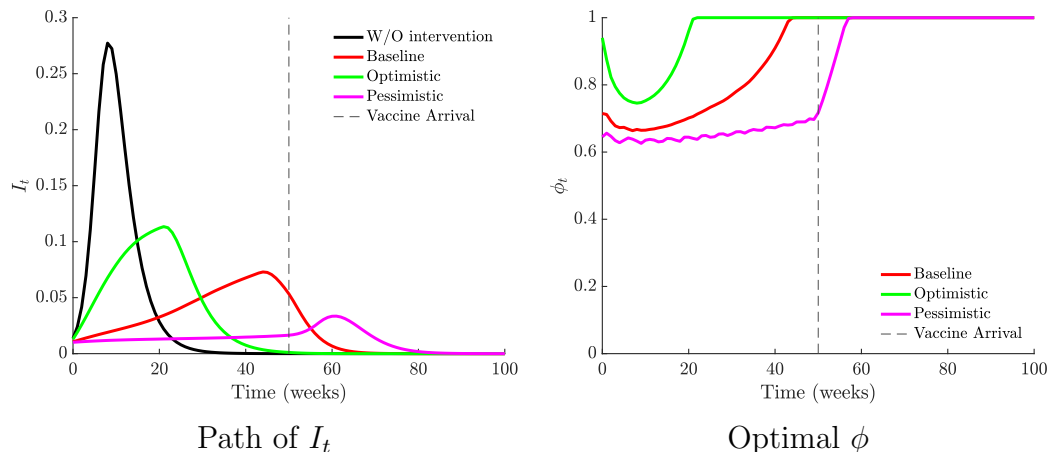


Figure 5. Optimistic and Pessimistic Scenarios

Relative to the baseline case, the implied paths for the epidemic and output are very different. The peak of the epidemic (in weeks since the start) can range from 22 to 60 weeks. In the first 3-4 weeks the loss of output is somewhat similar in both extreme cases. However, in the optimistic scenario recovery occurs early and it is very fast. Monthly output is almost 100% of normal after 7 months. On the other hand, in the pessimistic case it takes almost 1 1/2 year for output to return to normal.

Table 1 presents some summary statistics for the baseline and the two alternative scenarios. In addition to the relevant economic variables it shows the number of deaths averted and the output cost per death averted.

Table 1: Scenario Comparison			
Indicator	<i>Baseline</i>	<i>Optimistic</i>	<i>Pessimistic</i>
Output Loss (1 year)	22%	9.0%	35%
Output Loss (3 years)	7%	3.0%	12%
Full Recovery (months)	11	5.5	14.5
Deaths Averted	0.10%	0.04%	0.39%
Cost per Death Averted (Million \$)	12.6	12.8	5.5

These numbers illustrate sizeable differences for the optimistic and pessimistic cases relative to the baseline. The range is large both in terms of deaths averted (roughly a factor of 10) and the unit cost. These results comprise changes of multiple parameters. It is important to highlight how sensitive is the optimal response and the implied path of the epidemic to key parameters.

The impact of the fatality rate: The parameter χ controls the outflow of infected into deaths, an important component of the pay-off function. We explore the predictions of the model for three levels of the fatality rate: 0.375%, 0.75% and 1.13% which in terms of parameter

values correspond to setting χ equal to 0.025, 0.05 and 0.075. Figure 6 illustrate the effects of changes in the rate of fatalities.

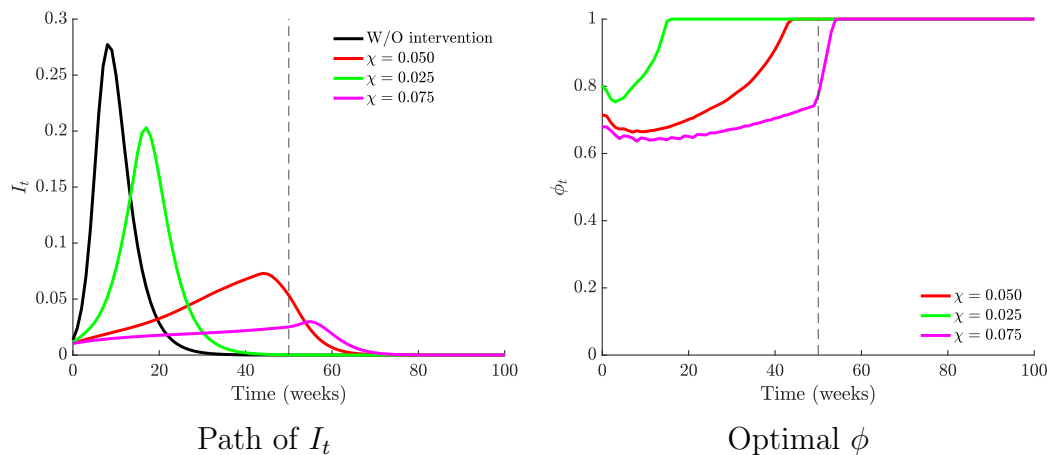


Figure 6. The Impact of the Fatality Rate

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 increase.

The impact of the speed of vaccination: The parameter $\bar{\mu}$ controls the speed of vaccination in Phase II, changing the payoff in this regime. We consider three maximum vaccination levels in terms of the time required to vaccinate 95% of the population: slow ($\bar{\mu} = 0.05$ which implies that it takes 60 weeks to attain that level of vaccination), medium (baseline, $\bar{\mu} = 0.10$, and 30 weeks), and fast ($\bar{\mu} = 0.15$, and 30 weeks). Figure 7 displays the time paths of infections and the optimal policy for these three maximum vaccination levels.




Figure 7. The Impact of Speed of Vaccination

The message is clear: the higher a country's ability to vaccinate its population the less severe the lockdown policy even before the vaccine becomes available. The reason for this is that when deciding on how much output to sacrifice the planner balances this cost with the potential cost of flattening the curve (postponing deaths to a later time). However, the faster the ability to vaccinate the lower the cost of postponing. Thus, to the extent that this variable captures some

dimension of the health infrastructure it implies that countries with poorer health infrastructure should adopt more restrictive policies.

The Duration of the Infectious Period: The parameter κ controls the duration of the infectiousness period. Relative to the baseline case that sets $\kappa = 1/3$ to capture 3 weeks, here we also consider the case of $\kappa = 1/2$ for a shorter infectiousness period of two weeks. depicted in Figure 8 with the green curve.




Figure 8. The Impact of Speed of Vaccination

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. This, in turn, implies that the optimal response is to apply a more severe lockdown policy in order to distribute potential deaths over time.²⁰

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 The Value of Life and the Cost of Averting Death

In our base formulation for the social value of a life, $\Delta(D) = M_0 D$, where M_0 corresponds to the utility equivalent of the value of a statistical life. In this section we present the results of three different values of M_0 . The value that we use in our baseline assumes that the present value of consumption of an individual who dies from the virus is \$347,000. We also consider

²⁰Even though the lockdown in the low κ case is more severe (lower ϕ) the speed at which the epidemic evolves is higher: The ratio of R_t^κ paths ($R_t^{0.5}/R_t^{0.33}$) is significantly greater than one. Thus, the planner does not fully compensate for the faster movement of the epidemic.

several other cases indexed by the equivalent present value of consumption (labeled PV in Table 2).

PV (\$)	Deaths Averted	Cost (Millions)	Y Loss (1Y) (%)	Trough (months)
1,330,000	0.35%	5.47	32	8
440,000	0.17%	8.87	27	3
347,000	0.10%	12.6	22	2
243,000	0.017%	19.6	5.9	3/4

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 foregone consumption as the metric). For the high valuations the number of deaths averted is significant, there is a very high output loss in the first year and the recovery does not start until after 4 months. In the low valuation case, the economy does not suffer much, very few deaths are averted at a very high unit cost, but the recession is short lived: after three months output starts increasing.

4.4 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 that 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. Thus, 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.

Table 3 reports the dollar equivalent of the gains associated with the availability of a vaccine for the base scenario and the optimistic and pessimistic cases described in Section 4.2.

Scenarios	Arrival Time (weeks)			
	1	4	25	50
Baseline	3.44	3.34	2.02	0.16
Optimistic	3.15	2.79	0.33	0.002
Pessimistic	3.07	3.03	2.56	1.91

We include the gains of having access to a vaccine at four points in time: at the beginning of the epidemic, a month into the epidemic, 6 months into the epidemic and a year after the beginning of the epidemic. Table 4 contains the same information except that it is indexed by

the present value of consumption as described in Section 4.3.

PV (\$)	Arrival Time			
	1	4	25	50
1,330,000	4.16	4.24	3.85	2.21
440,000	3.74	3.72	2.6	0.56
347,000	3.44	3.34	2.02	0.16
243,000	1.75	1.43	0.02	small

We view these results as indicative. On the one hand they overstate the gains by assuming that the alternative is no vaccine ever. On the other hand they ignore the value associated with a recurrence.²¹ The dollar values are large and there is a wide range of estimates. The main results include:

1. The initial value and the rate of depreciation do not move in sync. For example, when the infection rate is 1% the vaccine is more valuable in the optimistic case than in the pessimistic case but the ranking is reversed after four weeks (Table 3).
2. Lower vaccination capacity and higher fatality rate (Pessimistic case) and very high value of life (PV = 1,330) make the value of the vaccine higher and more persistent. The reason is that in both cases—for very different reasons—the planner finds it optimal to significantly flatten the curve and this implies that the epidemic does not extinguish itself rapidly and, hence, makes a vaccine valuable even after a year.
3. Not surprisingly, the market value of a vaccine is a monotone function of the value of a statistical life. This is driven by the fact that, in the model, a vaccine is the only technology that can lower the fatality rate.
4. Depending on the case a vaccine is significantly less valuable after 6 months and it has very little value after a year. To the extent that the private value (that is, the value of a patent) moves in sync with the social value, we would expect profit maximizing firms to decrease the investment in finding a vaccine as time goes by. Providing incentives using prizes rather than patents would eliminate the downward trend in investment.²²

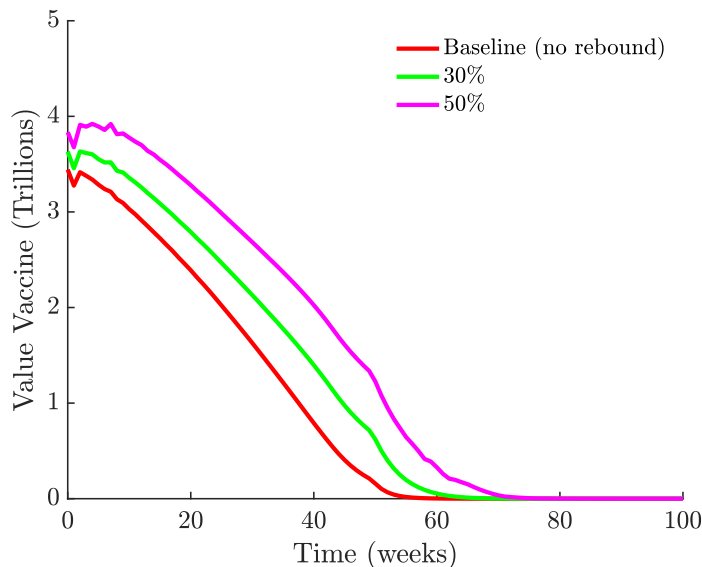
Recurrent Epidemics Our analysis of the value of a vaccine assumes that individuals who have been infected (and survived) have immunity against the cause of the epidemic. If this

²¹We discuss how to consider this in the section on extensions.

²²Of course, there is nothing suboptimal about the lower investment in the model as the social value of the vaccine is decreasing.

is not the case, then there is a possibility of a second wave. It is clear that the possibility of a second wave can only increase the value of producing a vaccine. The key question is by how much and, related, to what extent this possibility changes our finding that the value of a vaccine decreases significantly as time goes by.

To study the impact of a potential recurrence we assumed that recurrence is associated with a certain fraction of those who have immunity losing it. The arrival of this shock is modeled as a Poisson process. In Figure 9 we show the time path of the value of a vaccine in three cases: the baseline, one case in which 30% of the immune population instantaneously loses immunity and the second case when 50% losses immunity. The assumed arrival rate is $1/25$ which implies that a second wave would occur on average in approximately six months.²³



As expected the higher the size of the recurrence the higher the value of the vaccine. However, the normal path of the epidemic dominates and the value of a vaccine that is available after one year is a small fraction of its initial value.

5 Extensions

In this section we describe some extensions that are part of our ongoing work in this area.

Development The model can be used to ascertain how Phase I policies should be chosen depending on the ability to vaccinate the population rapidly. We are experimenting with a version of the model in which the fatality rate depends on whether infected individuals can have access to hospital facilities. Our preliminary results suggest that poor healthcare facilities

²³Assuming that recurrence occurs after 6 months rather than a year (which is more standard in the case of flu-like viruses) increases the value of the vaccine.

—as proxied by a low number of beds relative to demand— and low feasible vaccination rate —as captured by a low $\bar{\mu}$ — imply that a country with a more precarious public health infrastructure (for example a very poor country) should have a more aggressive policy —conditional on the state (S, I) — compared with a country with a good health sector.

In addition, we will move to a CRS utility function since the log function implies that income and substitution effects cancel.

Multi-Region Model We are extending the model to capture regional differences. We assume that each region produces two goods: one of the goods is produced with local inputs and the other uses inputs from other regions. In the first stage, we study the no-migration case and hence, in the epidemiological dimension we have N independent SIR models. We view the planner as choosing stay-at-home policies and vaccination rates (when available) for each region separately. We also compute the one-size-fits-all solution and study the welfare losses associated with that.

6 Concluding Comments

Here we offer some tentative comments about what we have learned from the exercise (in addition to the large impact that model uncertainty has on our predictions):

1. Stylized features of optimal policies.

- (a) Optimal stay-at-home policies policy depends on both the fraction infected and the fraction susceptible.
- (b) They have a shock treatment aspect to them: initially they are all severe but the duration depends on details of the economy.
- (c) The liberalization starts **before** the epidemic reaches its peak, and this sometimes results in an **increase** in the rate of infection.

2. Stylized features of suboptimal policies.

- (a) Policies that posit that liberalization must start after the epidemic peaks are suboptimal.
- (b) Policies that uniformly respond to increase in the rate of infection by tightening stay-at-home rules are suboptimal.

3. Vaccines.

- (a) Optimal pre-vaccine policies are not independent of the probability distribution of the arrival of a vaccine and the feasible vaccination rate.

- (b) The market value of a vaccine decreases rapidly (especially if the infection curve cannot be flattened).

4. **The Value of Life.**

- (a) The particular details how society puts a value on averting deaths has a first order effect on the optimal stay-at-home policy. This, in turn, determines the severity of the recession associated with the epidemic and the duration.
- (b) In many of our scenarios, the ex-post cost of averting a death is large and in the baseline case exceeds 5 million.

5. **The Role of Uncertainty.** 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] Acemoglu, D., V. Chernozhukov, I. Werning and M. Whinston, (2020), “A Multi-Risk SIR Model with Optimally Targeted Lockdown,” mimeo.
- [2] Alvarez, F., D. Argente and F. Lippi, (2020), “A Simple Planning Problem for COVID-19 Lockdown.” NBER working paper No 27102.
- [3] 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.
- [4] 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.
- [5] 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..
- [6] Baron, D, (2020), “Serology is Harder than It Looks,” Blog. April 14.
- [7] 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.
- [8] 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.
- [9] Bassetto, M. (2020), “Latent Covid-19 Diffusion in Italy: An Estimate from Italian Provinces,” Mimeo.
- [10] Bairoliya, N. and A. Imrohoroglu (2020), “Macroeconomic Consequences of Stay-At- Home Policies During the COVID-19 Pandemic,” Mimeo.
- [11] 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.
- [12] 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.
- [13] Bick, A. and A. Blandin (2020), “Real Time Labor Market Estimates During the 2020 Coronavirus Outbreak,” Mimeo.
- [14] Caballero, R. and A. Simpssek (2020), “ A Model of Asset Price Spirals and Aggregate Demand Amplification of a Covid-19 Shock,” Working Paper MIT.

- [15] Chang, R. and A. Velasco (2020), “Economic Policy Incentives to Preserve Lives and Livelihoods,” Working Paper.
- [16] Chari, V.V. and C. Phelan, (2020), “Test the Healthy and Gather Data we Don’t Already Have,” Minneapolis Star and Tribune (April 3)
- [17] Chen, X. and Z. Qiu, (2020), “Scenario Analysis of Non-Pharmaceutical Interventions on Global COVID-19 Transmissions,” working paper.
- [18] Correia, S., S. Luck, and E. Verner (2020), “ Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu,” Working Paper.
- [19] Eichenbaum, M.S., S. Rebelo, and M. Trabandt (2020a), “The Macroeconomics of Epidemics,” NBER Working Paper No 26882.
- [20] Eichenbaum, M.S., S. Rebelo, and M. Trabandt (2020b), “The Macroeconomics of Testing During Epidemics,” NBER Working Paper No 27104.
- [21] Dingel, J.I., and B. Neiman (2020), “How Many Jobs Can Be Done at Home?,” Working Paper Becker Friedman Institute.
- [22] Droz, L.A. and M.M. Tavares (2020), “Responding to COVID-19: A Note,” Working Paper Federal Reserve Bank of Philadelphia WP 20-14.
- [23] Farboodi, Maryam, Gregor Jarosch, and Robert Shimer, “Internal and External Effects of Social Distancing in a Pandemic,” NBER Working Paper No 27059.
- [24] 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.
- [25] Faria-e-Castro, (2020), “Fiscal Policy during a Pandemic,” Working Paper Federal Reserve Bank of St. Louis 2020-006A,
- [26] Ferguson, N. et. al. (2020), “Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand”
- [27] 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. Whitaker, 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.
- [28] 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.

- [29] Flaxman et. al., (2020), “Estimating the Number of Infections and the Impact of Non-Pharmaceutical Interventions on COVID-19 in 11 European Countries.”
- [30] Fornaro, L. and M. Wolf (2020), “Covid-19 Coronavirus and Macroeconomic Policy: Some Analytical Notes,” Working Paper CREI.
- [31] 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.
- [32] Gonzalez-Eiras and Niepelt (2020), “On the Optimal ‘Lockdown’ During an Epidemic” mimeo
- [33] Greenstone, M. and V. Nigam (2020), “Does Social Distancing Matter?,” Working Paper Becker Friedman Institute No 2020-26.
- [34] Gourinchas, P.O. (2020), “Flattening the Pandemic and Recession Curves,” Working Paper.
- [35] Gormsen, N.J. and R. S.J. Koijen (2020), “Coronavirus: Impact on Stock Prices and Growth Expectations,” Working Paper.
- [36] Guerrieri, V., G. Lorenzoni, L. Straub, and I. Werning (2020), “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?,” NBER
- [37] Hall, R.E., C.I. Jones, and P. J. Klenow (2020), “Trading-Off Consumption and COVID-19 Deaths,” Working Paper Stanford University.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] Kapicka, M. and P. Rupert (2020), “Labor Markets during Pandemics,” Working Paper Laboratory for Aggregate and Economics Fluctuations (LAEF).
- [42] Korolev, I, (2020), “Identification and Estimation of the SEIRD Epidemic Model for COVID-19,” working paper

- [43] Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020), “Scarring Body and Mind: The Long-Term Belief-Scarring Effects of COVID-19,” Working Paper.
- [44] Kubinec, R. (2020), “A Retrospective Bayesian Model for Measuring Covariate Effects on Observed COVID-19 Test and Case Counts,” Working Paper.
- [45] 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.
- [46] Krüger, D. H. Uhlig, and T. Xie (2020), “Macroeconomic Dynamics and Reallocation in an Epidemic,” Working Paper Becker Friedman Institute No 2020-43.
- [47] Manski, C. and F. Molinari (2020), “Estimating the COVID-19 Infection Rate: Anatomy of an Inference Problem,” Working Paper.
- [48] 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.
- [49] 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).
- [50] Pindyck R.S (2020), “COVID-19 and the Welfare Effects of Reducing Contagion,” NBER Working Paper No. 27121.
- [51] Shao, P. (2020), “Impact of city and residential unit lockdowns on prevention and control of COVID-19,” Working Paper.
- [52] Stock, J.H. (2020), “Data Gaps and the Policy Response to the Novel Coronavirus James H. Stock,” NBER Working Paper No. 26902.
- [53] 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.
- [54] 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.

Appendix 1: Proofs

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

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

The associated Hamiltonian is

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

The FOC are standard and given by the static conditions

$$\begin{aligned} u^\theta()w(1 - \zeta I) &= 2\beta\phi(1 - \zeta)SI(\lambda_S - \lambda_I) + \hat{\gamma}^\phi, \\ \hat{\gamma}^\phi(1 - \phi) &= 0 \end{aligned} \tag{8}$$

and

$$-u^\theta()c_V^\theta(\mu(S + (1 - \zeta)I))(S + (1 - \zeta)I) = -\lambda_S S + \hat{\gamma}^I - \hat{\gamma}_+^I, \tag{9}$$

where we omit the argument 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 $\phi \in (0, 1)$, equation (8) can be written as, given that $\lambda_S = F_S$ and $\lambda_I = F_I$,

$$\frac{u^\theta(\phi w(1 - \zeta I) - c_V(\mu(S + (1 - \zeta)I)))w(1 - \zeta I)}{2\beta\phi(1 - \zeta)SI} = (F_S - F_I),$$

which corresponds to equation (7) in the text.

Proof. Consider first Phase II. The relevant co-state variables evolve according to the following differential equations

$$\begin{aligned} \dot{\lambda}_S = & u^\theta(\phi w(1 - \zeta I) - c_V(\mu(S + (1 - \zeta)I)))c_V^\theta(\mu(S + (1 - \zeta)I))\mu \\ & + (\rho + \mu + \gamma)\lambda_S + (\lambda_S - \lambda_I)\beta\phi^2(1 - \zeta)I \end{aligned} \tag{10}$$

$$\begin{aligned} \dot{\lambda}_I = & u^\theta(\phi w(1 - \zeta I) - c_V(\mu(S + (1 - \zeta)I))) [w\phi\zeta + c_V^\theta(\mu(S + (1 - \zeta)I))(1 - \zeta)\mu] \\ & + \Delta^\theta[(\chi \zeta I - \bar{H})^+]\chi\zeta + (\rho + \kappa)\lambda_I + (\lambda_S - \lambda_I)\beta\phi^2(1 - \zeta)S + \lambda_S\gamma. \end{aligned} \tag{11}$$

The static first order conditions are equations (8) and (9).

We conjecture that there is a steady state such that $\phi = 1$ and $\mu = 0$. At this steady state the epidemiological variables satisfy

$$\beta(1 - \zeta)S = \kappa, \quad (12)$$

and

$$I = \frac{\gamma}{\gamma + \kappa}(1 - S). \quad (13)$$

It suffices to show that the system of equations that is implied by $\dot{\lambda}_S = \dot{\lambda}_I = 0$ has a solution evaluated at the candidate steady state and that satisfies equations (8) and (9).

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

$$\lambda_S = -\frac{\beta(1 - \zeta)I u^\theta(w(1 - \zeta I))\zeta w}{\Lambda} \quad (14)$$

and

$$\lambda_I = -\frac{(\gamma + \rho + \beta(1 - \zeta)I) u^\theta(w(1 - \zeta I))\zeta w}{\Lambda}, \quad (15)$$

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 (8) and (9) hold as inequalities (ignoring the Lagrange multipliers). Some standard manipulations show that this is equivalent to (in the case of equation (8)) to

$$(1 - \zeta I) > 2\kappa I \left[\frac{(\rho + \gamma)\zeta}{\Lambda} \right].$$

Since the left hand side is decreasing in I and it converges to 1 as $I \rightarrow 0$, while the right hand side converges to zero as $I \rightarrow 0$ then equation (8) is satisfied. To check that equation (9) holds as well, it suffices to show that

$$u^\theta(w(1 - \zeta I))c_V^\theta(0) \geq -\lambda_S \frac{S}{S + I} = \frac{S}{S + I} \frac{\beta(1 - \zeta)I u^\theta(w(1 - \zeta I))\zeta w}{\Lambda},$$

and this holds for I sufficiently small since $\lim_{I \rightarrow 0} \Lambda = \rho(\rho + \gamma) > 0$.

Now we want to show that 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

$$\begin{aligned} H^I = & u(\phi w L) - \Delta [(\chi \zeta I_t - \bar{H})^+] + \eta F(S, I) - \pi_S [\beta \phi^2 (1 - \zeta) S I + \gamma (1 - S - I)] \\ & + \pi_I [\beta \phi^2 (1 - \zeta) S I - \kappa I], \end{aligned}$$

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

We assume that the function $F(S, I)$ is differentiable (to be proved later) 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 the limiting $\phi = 1$. The steady state is such that

$$\pi_S = \frac{\eta(\rho + \eta)F_S + \beta(1 - \zeta)I(\eta F_S - u^\theta(w(1 - \zeta I))w\zeta)}{\tilde{\Lambda}},$$

$$\pi_I = \frac{(\rho + \eta + \gamma + \beta(1 - \zeta)I) (\eta F_S - u^0(w(1 - \zeta I))w\zeta) + \eta F_S (\beta(1 - \zeta)I)}{\tilde{\Lambda}},$$

where

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

and

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

It follows that $\lim_{I \rightarrow 0} \pi_S = 0$ (details omitted but just brute force) and $\lim_{I \rightarrow 0} \pi_I < 0$ (and finite)

The relevant first order condition to guarantee that the solution to the static condition is $\phi = 1$ is

$$(1 - \zeta I) > 2\kappa I [\pi_S - \pi_I],$$

and it is clearly satisfied for small I .

To summarize if

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

is sufficiently small (that is, if γ —the rate at which the population of susceptibles is replenish) is small) then the long run behavior with and without vaccines is exactly the same. ■

Appendix 2: Computational Appendix

Writing the HJB equations for the two phases,

$$\begin{aligned}\rho F(S, I) &= \max_{(\phi, \bar{1})} \{u(\phi w(1 - \zeta I) - c_V(\mu(S + (1 - \zeta)I))) - \Delta[(\chi \zeta I_t - \bar{H})^+]\} \\ &\quad + F_S[-\beta \phi^2(1 - \zeta)SI - \mu S + \gamma(1 - S - I)] + F_I[\beta \phi^2(1 - \zeta)SI - \kappa I]\} \\ \rho V(S, I) &= \max_{(\phi, \bar{1})} \{u(\phi w(1 - \zeta I) - c_V(\mu(S + (1 - \zeta)I))) - \Delta[(\chi \zeta I_t - \bar{H})^+]\} \\ &\quad + V_S[-\beta \phi^2(1 - \zeta)SI + \gamma(1 - S - I)] + V_I[\beta \phi^2(1 - \zeta)SI - \kappa I]\} + \eta(F(S, I) - V(S, I))\end{aligned}$$

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 non-uniform grid and restrict the space to $0 \leq S + I \leq 1$.

Phase II

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

FOC:

$$\begin{aligned}\frac{(1 - e^{-\rho \Delta t})}{\rho} \frac{(1 - \zeta I)}{\phi(1 - \zeta I) - (1 - c_0)} &= e^{-\rho \Delta t} [F_S(S_{t+}, I_{t+}) - F_I(S_{t+}, I_{t+})] [2\beta \phi(1 - \zeta)SI] \Delta \\ \phi[\phi(1 - \zeta I) - (1 - c_0)] &= \frac{(1 - e^{-\rho \Delta t})}{\Delta \rho e^{-\rho \Delta t}} \frac{(1 - \zeta I)}{[2\beta(1 - \zeta)SI] [F_S(S_{t+}, I_{t+}) - F_I(S_{t+}, I_{t+})]} \\ \phi^2(1 - \zeta I) - \phi(1 - c_0) - \frac{(1 - e^{-\rho \Delta t})}{\Delta \rho e^{-\rho \Delta t}} &\frac{(1 - \zeta I)}{[2\beta(1 - \zeta)SI] [F_S(S_{t+}, I_{t+}) - F_I(S_{t+}, I_{t+})]} = 0\end{aligned}$$

Define:

$$d \equiv \frac{(1 - e^{-\rho \Delta t})}{\Delta \rho e^{-\rho \Delta t}} \frac{1}{[2\beta(1 - \zeta)SI] [F_S(S_{t+}, I_{t+}) - F_I(S_{t+}, I_{t+})]}$$

$$\begin{aligned}\phi^2(1 - \zeta I) - \phi(1 - c_0) - d(1 - \zeta I) &= 0 \\ \phi &= \frac{(1 - c_0) \pm \sqrt{(1 - c_0)^2 + 4(1 - \zeta I)^2 d}}{2(1 - \zeta I)}\end{aligned}$$

Given that we want $c \geq (1 - c_0)w$, we can ignore the lower root. Thus,

$$\phi = \frac{(1 - c_0) + \sqrt{(1 - c_0)^2 + 4(1 - \zeta I)^2 d}}{2(1 - \zeta I)}$$

Phase I

$$\mu = 0$$

$$V(S_t, I_t) = \max_{(0 \leq \phi \leq 1)} \left\{ \frac{(1 - e^{-(\rho+\eta)})}{\rho + \eta} \left(u(\phi w(1 - \zeta I) - c_V(\mu S)) - \Delta [D_t] + \eta F(S_t, I_t) \right) + e^{-(\rho+\eta)} V(S_{t+}, I_{t+}) \right\}$$

$$S_{t+} = S_t + [-\beta \phi^2 (1 - \zeta) S I - \mu S + \gamma(1 - S - I)] \Delta$$

$$I_{t+} = I_t + [\beta \phi^2 (1 - \zeta) S I - \kappa I] \Delta$$

FOC:

$$\phi = \frac{(1 - c_0) + \sqrt{(1 - c_0)^2 + 4(1 - \zeta I)^2 d}}{2(1 - \zeta I)}$$

Where,

$$d \equiv \frac{(1 - e^{-(\rho+\eta)})}{\Delta(\rho + \eta)e^{-(\rho+\eta)}} \frac{1}{[2\beta(1 - \zeta) S I] [V_S(S_{t+}, I_{t+}) - V_I(S_{t+}, I_{t+})]}$$

Appendix 3: Calibration

1. **Utility.** We consider log utility. To be precise we assume that

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

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

We assume that $c_0 = 0.4$, and, hence, that output cannot fall below 60% of its steady state value

2. **Vaccination.**

- (a) The *cost of vaccination*:

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

The value of c_V^0 depends on the units in the function u .

Baseline: $c_V^0 = 0$.

- (b) *Vaccination capacity.* We set $\bar{\mu} = 0.10$

3. **Epidemiological Parameters.**

- (a) The *Fraction* ζ . A difficulty estimating ζ is the lack of random testing at this point and the as-hoc assumptions about mortality that have to be made to produce estimates. Hortacsu et. al. 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) using a different approach and relying on Chinese data estimate $\zeta = 0.04$. More recent evidence give a range between 0.10 and 0.30. We assume $\zeta = 0.16$

- (b) *The Fatality Rate.* What are the implied case fatality rates? Given that the case fatality rate is $\chi\zeta I$, the true fatality rate is $\chi\zeta$. Thus, we need an estimate of χ which corresponds to the ratio of fatalities/diagnosed cases. In the US at the time of this writing, the measured death rate is about 6%, while for the world as a whole it exceeds 6%. At the same time there some countries —including several Latin American countries— in which the case fatality rate is below 4%. Thus, it seems that $\chi = 0.05$ is a reasonable estimate.

The implied fatality rate is $0.05 \times \zeta$. If we assume that $\zeta = 0.15$ the fatality rate is 0.75%, which is in the range of estimates. The lower bound of the estimates of ζ (around 0.04) implied a fatality rate of 0.2% which is slightly higher than the influenza fatality rate. The upper bound of the estimates ($\zeta = 0.2$) implies a true fatality rate equal to 1%.

Baseline: $\zeta = 0.15$

- (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

$$\beta = \frac{\kappa \times \mathcal{R}_0}{1 - \zeta} = \frac{0.16 \times 2.8}{0.9} = 0.497.$$

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: $\beta = 0.497$

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* . Since the constant is given by

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

If we assume that $v = 1$ the implied present value of consumption (using the same life expectancy) is 450,000 which is on the high side. In the baseline we assume that the present discounted value of consumption is 350,000. Thus we have the following values.

Present Value of Consumption (\$)	M_0
1,330,000	3,797
440,000	2,994
347,000	2,693
243,000	1,807

(b) **Definition:** \mathcal{R}_t

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

6. Next we pick γ . Fix the steady state I , then (modulo algebraic error) I get that

$$\gamma = \frac{\kappa I}{(1 - I - \frac{1}{R_0})}.$$

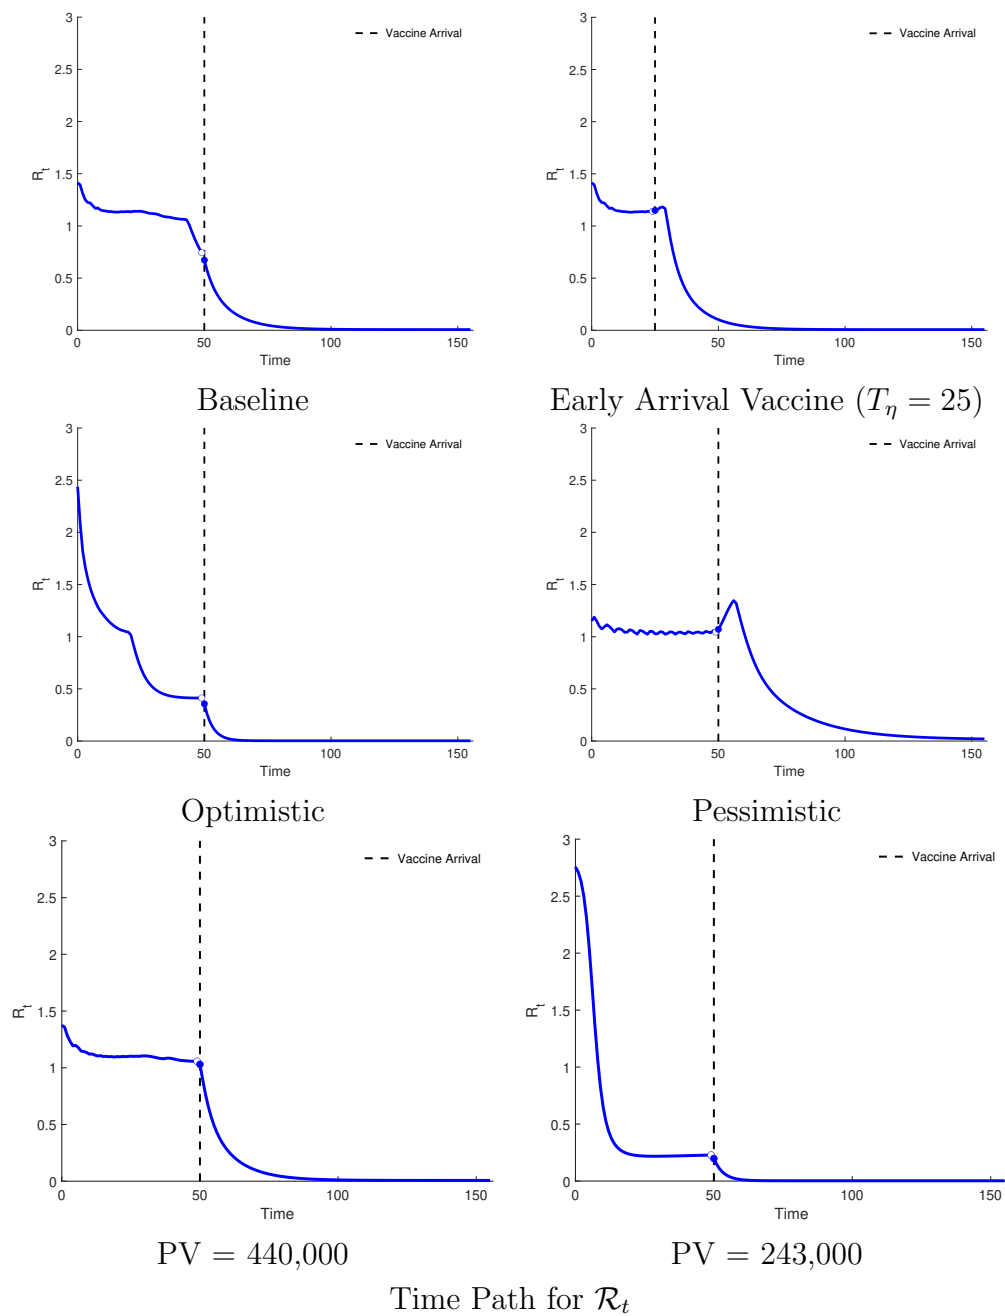
If $I = 0.001$ then I get that $\gamma = .00025$.

Table 1: Baseline Parameters


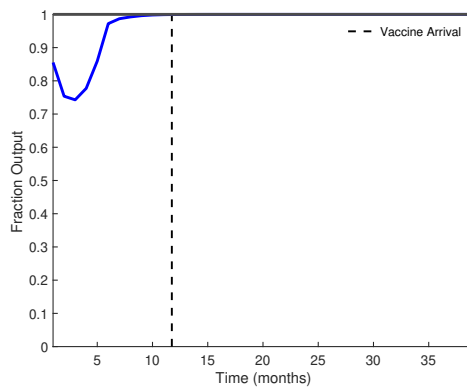

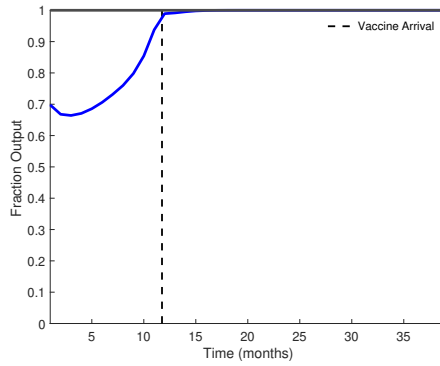
Meaning	Parameter	Value
Fraction Diagnosed among Infected	ζ	0.15
Immunity Loss Rate	γ	0.00025
Basic Reproduction Number	R_0	2.8
Recovery Rate	κ	0.16
Discount Rate	ρ	0.000233
Time Step	Δ	1/15
Loss function	M_0	2693
Output per Worker	w	1173
Case Fatality Rate	χ	0.05
Hospital Capacity	\bar{H}	0.000427
Vaccine Cost	c_v	0
Vaccination capacity	μ	0.10
Minimum Consumption (1- c_0)	c_0	0.60
Vaccine Arrival Poisson	η	1/50
Initial I_0	I_0	0.01
Initial S_0	S_0	0.99

Appendix 4: Supplementary Graphs

- Here we present the model's implications for measured \mathcal{R}_t in several cases.



- Time series for monthly output



PV = 440,000

PV = 243,000

Monthly Output (full employment = 100)

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 difference 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 and Niepelt (2020), Acemoglu, Chernozhukov, Werning, and Whinston (2020), Jones, Philippon, and Venkateswaran (2020).

Policies in an SIR models: These papers explore the effects of different policies in the dynamics of the pandemic (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 Schweg 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 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 Imrohoroglu (2020), Krüger, Uhlig, and Xie (2020), Kozlowski, Veldkamp, Venkateswaran (2020), Bodenstein, Corsetti, and Guerrieri (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).