

A Note on Efficient Mitigation Policies*

Callum Jones[†] Thomas Philippon[‡] Venky Venkateswaran[§]

First version: March 31, 2020; this version: April 7, 2020

Abstract

We study the response of an economy to an unexpected epidemic and we compare the decentralized equilibrium with the efficient allocation. Households mitigate the spread of the disease by reducing consumption and hours worked. A social planner worries about two externalities: an infection externality and a healthcare congestion externality. Private agents' mitigation incentives are too weak, especially at early stages while the planner implements drastic and front-loaded mitigation policies. In our calibration, assuming a CFR of 1% and an initial infection rate of 0.1%, private mitigation leads to a 10% drop in consumption and reduces the cumulative death rate from 2.5% of the initially susceptible population to about 2%. The planner reduces the death rate to 0.2% at the cost of an initial drop in consumption of around 40%.

Keywords: contagion, containment, covid 19, recession, R_0 , social distancing, *SIR* model, mitigation, suppression, vaccine.

1 Introduction

The response to the Covid-19 crisis highlights the tension between health and economic outcomes. The containment measures that can help slow the spread of the virus are likely to reinforce the economic downturn. Policy makers have naturally recognized this trade off and we hope to contribute to ongoing effort to provide quantitative models to guide their decisions.

*We are grateful for comments from seminar participants at the University of Chicago. The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, or its management. First version: March 31, 2020.

[†]International Monetary Fund, jonescallum@gmail.com

[‡]New York University, CEPR and NBER, tphilipp@stern.nyu.edu

[§]New York University, NBER, vvenkate@stern.nyu.edu

We propose a simple extension of the neoclassical model to quantify the tradeoffs and guide policy. We are particularly interested in understanding the design of the policy response. When will the private sector engineer the right response, and when is there a need for policy intervention? Which measures should be front-loaded and which ones should ramp up as the contagion progresses?

Our model has two building blocks: one for dynamics of contagion, and one for consumption and production. Our starting point is the classic *SIR* model of contagion used by public health specialists. Atkeson (2020b) provides a clear summary of this class of model. In a population of initial size N , the epidemiological state is given by the numbers of Susceptible (S), Infected (I), and Recovered (R) people. By definition, the cumulative number of deaths is $D = N - S - I - R$. Infected people transmit the virus to susceptible people at a rate that depends on the nature of the virus and on the frequency of social interactions. Containment, testing, and social distancing reduce this latter factor. The rates of recovery (transitions from I to R), morbidity (I becoming severely or critically sick) and mortality (transition from I to D) depend on the nature of the virus and on the quality of health care services. The quality of health services depends on the capacity of health care providers (ICU beds, ventilators) and the number of sick people.

On the economic side of the model we use a standard model where members of large households jointly make decisions about consumption and labor supply. We assume that the consumption of (some) goods and services increases the risk of contagion, and that going to work also increases the risk of contagion.

We can then study how the private sector reacts to the announcement of an outbreak and how a government should intervene. Upon learning of the risks posed by the virus, households change their labor supply and consumption patterns. They cut spending and labor supply in proportion to the risk of infection, which – all else equal – is proportional to the fraction of infected agents I/N . Households only take into account the risk that they become infected, not the risk that they infect others, therefore their mitigation efforts are lower than what would be socially optimal. This infection externality is well understood in the epidemiology literature. The other important externality is the congestion externality in the healthcare system. When hospitals are overwhelmed the risk of death increases but agents do not internalize their impact on the risk of others.

We obtain interesting results when we compare the timing of mitigation. The planner wants to front load these efforts compared to the private sector. The risk of future contagion and of congestion in the health care system also drives an important wedge between private decisions and the socially

efficient allocation. If a private agent knows that she is likely to be infected in the future, this reduces her incentives to be careful today. We call this effect the fatalism effect. The planner on the other hand, worries about future infections and future congestion.

Literature Our paper relates to the literature on contagion dynamics (Diekmann and Heesterbeek, 2000). We refer to the reader to Atkeson (2020b) for a recent discussion. Berger et al. (2020) show that testing can reduce the economic cost of mitigation policies as well as reduce the congestion in the health care system. Baker et al. (2020) document the early consumption response of US households.

The most closely related papers are Barro et al. (2020), Correia et al. (2020), Eichenbaum et al. (2020) and Alvarez et al. (2020). Barro et al. (2020) and Correia et al. (2020) draw lessons from the 1918 flu epidemic. Barro et al. (2020) find a high death rate (about 40 million people, 2% of the population at the time) and a large but not extreme impact on the economy (cumulative loss in GDP per capita of 6% over 3 years). The impact on the stock market was small. Correia et al. (2020) find that early interventions help protect health and economic outcomes.

Our model shares with Eichenbaum et al. (2020) and Alvarez et al. (2020) the idea of embedding *SIR* dynamics in a simple DSGE model. The *SIR* model is the same, but there are some key differences in the DSGE model. Eichenbaum et al. (2020) consider hand-to-mouth agents who know their true status (infected or susceptible) while we work within a representative shopper/worker household framework à la Lucas and Stokey (1987). We explain the dynamic tension between the planner and the private sector and we describe an endogenous fatalism bias in private incentives. In terms of results, we find that optimal interventions are more front-loaded in our setting than in Eichenbaum et al. (2020). In their benchmark case, they find that the optimal policy is to gradually ramp up containment policies as the number of infected people rises, while congestion externalities in the healthcare system induce more aggressive, albeit still gradual, containment. Front-loading seems to become optimal only when agents anticipate the arrival of a vaccine in the future. In our setting, front-loading is optimal in almost all cases.

Alvarez et al. (2020) study a lockdown planning problem under *SIR* dynamics. They assume risk neutral agents and a linear lockdown technology. They find that the congestion externality plays an important role in shaping the policy response and that the planner front-loads the effort. Our planner has similar incentives but takes into account the desire for consumption smoothing. Jones et al. (2020) study optimal policies in a pandemic, including the option of working from home.

2 Benchmark Model

2.1 Households

There is a continuum of mass N of households. Each household is of size 1 and the utility of the household is

$$U = \sum_{t=0}^{\infty} \beta^t u(c_t, l_t; i_t, d_t),$$

where c_t is per-capita consumption and l_t is labor supplied by those who are alive and not sick. The household starts with a continuum of mass 1 of family members, all of them susceptible to the disease. At any time $t > 0$ we denote by s_t , i_t and d_t the numbers of susceptible, infected and dead people. The size of the household at time t is therefore $1 - d_t$. If per capita consumption is c_t then household consumption is $(1 - d_t) c_t$. Among the i_t infected members, κi_t are too sick to work. The labor force at time t is therefore $1 - d_t - \kappa i_t$, and household labor supply is $(1 - d_t - \kappa i_t) l_t$. The number of household members who have recovered from the disease is $r_t = 1 - s_t - i_t - d_t$. In the quantitative applications below we use the functional form

$$u(c_t, l_t; i_t, d_t) = (1 - d_t - \kappa i_t) \left(\log(c_t) - \frac{l_t^{1+\eta}}{1+\eta} \right) + \kappa i_t (\log(c_t) - u_\kappa) - u_d d_t,$$

where u_κ is the disutility from being sick and u_d the disutility from death which includes lost consumption and the psychological cost on surviving members.¹ For simplicity we assume that sickness does not change the marginal utility of consumption therefore c is the same for all alive members of the household. The variables s , i and d evolve according a standard *SIR* model described below. We use a Lucas and Stokey (1987) approach to model households. At the beginning of period t the household decides how much to consume c_t (per capita) and how much each able-bodied member should work l_t . Then the shoppers go shopping and the workers go to work. Notice that we have normalized the disutility of labor so that $l = c = 1$ before the epidemic starts.

Households understand that they can become infected by shopping and by going to work. We compute infection in two steps. First we define exposure levels for shoppers and for workers. Then we aggregate these into one infection rate at the household level. Finally we take into account the stochastic arrival of a vaccine by adjusting the discount factor β . Formally, we assume an exogenous arrival rate for a cure to the disease. By a cure we mean both a vaccine and a treatment for the

¹Formally $u_d = PsyCost - \log(c_d)$ where c_d is the consumption equivalent in death. Technically we cannot set $c_d = 0$ with log preferences but u_d is a large number.

currently sick. Under this simplifying assumption the economy jumps back to $l = c = 1$ when a cure is found. We can therefore focus on the stochastic path before a cure is found. Let $\tilde{\beta}$ be the pure time discount rate and ν the likelihood of a vaccine. We define $\beta = \tilde{\beta}(1 - \nu)$ along the no-cure path.

2.2 Shopping

Household members can get infected by shopping. We define consumption (shopping) exposure as

$$e^c c_t C_t,$$

where e^c measures the sensitivity of exposure to consumption and C_t is aggregate consumption, all relative to a steady state value normalized to one. The idea behind this equation is that household members go on shopping trips. We assume that shopping trips scale up with consumption and that, for a given level of aggregate consumption, exposure is proportional to shopping trips. This functional form captures the notion of crowds in shopping mall as well as in public transportation.

2.3 Production

Exposure at work for household members working is given by

$$e^l l_t L_t,$$

where e^l measure the sensitivity of exposure to labor and L_t is aggregate labor supply. Effective labor supply, \hat{l}_t , is given by

$$\hat{l}_t = (1 - d_t - \kappa i_t) l_t.$$

This equation captures the fact that the number of valid household member is decreased by death and sickness. Production is linear in effective labor

$$Y_t = \hat{L}_t = N \hat{l}_t.$$

In our basic model we ignore the issue of firm heterogeneity and market power. Therefore price is equal to marginal cost

$$P_t = W_t = 1,$$

where W is the wage per unit of effective labor, which we normalize to one.

2.4 Income and Contagion

At the end of each period, household members regroup and share income, consumption and exposure.

Household labor income is $W_t \hat{l}_t = \hat{l}_t$ and the budget constraint is

$$(1 - d_t) c_t + \frac{b_{t+1}}{1 + r_t} \leq b_t + \hat{l}_t$$

Household exposure is

$$e_t = \bar{e} + (1 - d_t) e^c c_t C_t + (1 - d_t - \kappa i_t) e^l l_t L_t,$$

where \bar{e} is baseline exposure, independent of market activities. Contagion dynamics follow an *SIR* model. From the perspective of one household, this model is:

$$\begin{aligned} s_{t+1} &= s_t - \gamma e_t \frac{I_t}{N} s_t \\ i_{t+1} &= \gamma e_t \frac{I_t}{N} s_t + (1 - \rho) i_t - \delta_t \kappa i_t \\ d_{t+1} &= d_t + \delta_t \kappa i_t \\ r_{t+1} &= r_t + \rho i_t \end{aligned}$$

where γ is the infection rate per unit of exposure, ρ the recovery rate, κ the probability of being sick conditional on infection, and δ_t the mortality rate of sick patients. In the standard *SIR* model γ is constant. In our model it depends on exposure and therefore on mitigation strategies. The parameter δ_t increases when the health system is overwhelmed, as discussed below.

2.5 Market Clearing and Aggregate Dynamics

Infection dynamics for the the entire population are simply given by the *SIR* system above with aggregate variable $I_t = N i_t$, and so on. The aggregate labor force is $N(1 - \kappa i_t - d_t) l_t$ and total consumption is $N(1 - d_t) c_t$. Goods market clearing requires

$$(1 - d_t) c_t = \hat{l}_t,$$

and bond market clearing requires

$$b_t = 0.$$

Finally we capture the limited capacity of the healthcare system with the increasing function

$$\delta_t = \delta(I_t).$$

Note that $\delta(I_t)$ should really be written as $\delta(\kappa I_t, H_t)$ where κI_t is the number of sick people and H_t is the capacity of the healthcare system. Since we assume that both κ and H are constant we write simply $\delta(I_t)$. We call the fact that δ is increasing the congestion externality.

3 Decentralized equilibrium

Our main goal is to compare the decentralized equilibrium with the planner's solution.

3.1 Equilibrium Conditions

Since our model reduces to a representative household model and since $b = 0$ in equilibrium, we simply omit b from the value function. The household's recursive problem is

$$V_t(i_t, s_t, d_t) = \max_{c_t, l_t, m_t} u(c_t, l_t; i_t, d_t) + \beta V_{t+1}(i_{t+1}, d_{t+1}, s_{t+1}),$$

where the flow utility is

$$u(c_t, l_t; i_t, d_t) = (1 - d_t) \log(c_t) - (1 - d_t - \kappa i_t) \frac{l_t^{1+\eta}}{1+\eta} - u_\kappa \kappa i_t - u_d d_t$$

Using the definition of effective labor $\hat{l}_t = (1 - d_t - \kappa i_t) l_t$, we can write the Lagrangian as

$$\begin{aligned} V_t = & u(c_t, l_t; i_t, d_t) + \beta V_{t+1} + \lambda_t \left(\hat{l}_t + b_t - (1 - d_t) c_t - \frac{b_{t+1}}{1 + r_t} \right) \\ & + \lambda_{e,t} \left(e_t - \bar{e} - (1 - d_t) e^c c_t C_t - (1 - d_t - \kappa i_t) e^l l_t L_t \right) \\ & + \lambda_{i,t} \left(i_{t+1} - \gamma e_t \frac{I_t}{N} s_t - (1 - \rho) i_t + \delta_t \kappa i_t \right) \\ & + \lambda_{s,t} \left(s_{t+1} - s_t + \gamma e_t \frac{I_t}{N} s_t \right) \\ & + \lambda_{d,t} (d_{t+1} - d_t - \delta_t \kappa i_t) \end{aligned}$$

We highlight in red the externalities, from infection and from congestion. The first order conditions for consumption and labor are then

$$c_t : c_t^{-1} = \lambda_t + \lambda_{e,t} e^c C_t$$

$$l_t : l_t^\eta = \lambda_t - \lambda_{e,t} e^l L_t$$

The remaining first order conditions are

$$e_t : \lambda_{e,t} = (\lambda_{i,t} - \lambda_{s,t}) \gamma \frac{I_t}{N} s_t$$

$$i_{t+1} : \lambda_{i,t} = -\beta V_{i,t+1}$$

$$s_{t+1} : \lambda_{s,t} = -\beta V_{s,t+1}$$

$$d_{t+1} : \lambda_{d,t} = -\beta V_{d,t+1}$$

The envelope conditions are

$$V_{i,t} = \kappa \frac{l_t^{1+\eta}}{1+\eta} - \kappa u_\kappa - \kappa \lambda_t l_t + \lambda_{e,t} \kappa e^l l_t L_t - (1-\rho) \lambda_{i,t} + \delta_t \kappa (\lambda_{i,t} - \lambda_{d,t})$$

$$V_{s,t} = (\lambda_{s,t} - \lambda_{i,t}) \gamma e_t \frac{I_t}{N} - \lambda_{s,t}$$

$$V_{d,t} = \frac{l_t^{1+\eta}}{1+\eta} - \log(c_t) - u_d - \lambda_t (l_t - c_t) + \lambda_{e,t} (e^c c_t C_t + e^l l_t L_t) - \lambda_{d,t}$$

3.2 Equilibrium with Exogenous Infections

Economy before the pandemic To simplify the notation we normalize $N = 1$, so we should think of our values as being per-capita pre-infection. When there is no risk of contagion, i.e., when $i_t = 0$ and $\lambda_{e,t} = 0$, optimal consumption and labor supply implies $c_t^{-1} = l_t^\eta$. We have $\hat{l}_t = l_t$ so market clearing is simply $c_t = l_t$. Combining these two conditions we get

$$c_t = l_t = 1.$$

The pre-infection economy is always in steady state.

Exogenous infections Consider now an economy with exogenous SIR dynamics: $e^c = e^l = 0$. The SIR system is then independent from the economic equilibrium. In the SIR system, the share of infected agents I_t increases, reaches a maximum and converges to 0 in the long run. Assuming a constant δ , the long run solution solves

$$\log\left(\frac{S_\infty}{1-I_0}\right) = -\frac{\gamma\bar{e}}{\rho+\delta\kappa}\left(\frac{1-S_\infty}{N}\right),$$

and

$$D_\infty = \frac{\delta\kappa}{\delta\kappa+\rho}(1-S_\infty).$$

When the congestion externality arises and δ_t increases, then we cannot obtain a closed-form solution for the long run death rate but the qualitative results are unchanged. Since $e^c = e^l = 0$ we have $m_t = 0$ and $c_t^{-1} = l_t^\eta$. Market clearing requires $(1-d_t)c_t = (1-d_t-\kappa i_t)l_t$ therefore labor supply is

$$l_t^{1+\eta} = 1 + \frac{\kappa i_t}{1-d_t-\kappa i_t}.$$

The labor supply of valid workers increases to compensate for the reduced productivity of the sick. Per capita consumption is

$$c_t = \left(\frac{1-d_t}{1-d_t-\kappa i_t}\right)^{-\frac{\eta}{1+\eta}}$$

As long as $\eta > 0$ consumption per capita decreases. Aggregate GDP decreases because of lost labor productivity and deaths. The following proposition summarizes our results.

Proposition 1. *When contagion does not depend on economic activity ($e^c = e^l = 0$), the share of infected agents I_t increases, reaches a maximum and converges to 0 in the long run. The long run death rate is given by $D_\infty = \frac{\delta\kappa}{\delta\kappa+\rho}(1-S_\infty)$ where the long run share of uninfected agents solves $\log\left(\frac{S_\infty}{1-I_0}\right) = -\frac{\gamma\bar{e}}{\rho+\delta\kappa}\left(\frac{1-S_\infty}{N}\right)$. Along the transition path, labor supply of able-bodied workers follows the infection rate while per-capita consumption moves in the opposite direction as $c_t = \left(1 - \frac{\kappa i_t}{1-d_t}\right)^{\frac{\eta}{1+\eta}}$.*

3.3 Private Incentives for Mitigation

Let us focus on consumption by setting $e^l = 0$. Optimal private consumption is

$$c_t^{-1} = \lambda_t + \lambda_{e,t}e^c C_t,$$

so the temptation to cut consumption depends on $\lambda_{e,t} = (\lambda_{i,t} - \lambda_{s,t}) \gamma \frac{I_t}{N} s_t$ which is high when $\gamma \frac{I_t}{N} s_t$ is high, which is exactly when new infections are high and S is quickly decreasing. So holding constant $\lambda_{i,t} - \lambda_{s,t}$ the private incentives to cut consumption are proportional to the number of new cases. The other important element is

$$\lambda_{i,t} - \lambda_{s,t} = \beta (V_{s,t+1} - V_{i,t+1})$$

The right-hand-side of this expression represents the value of avoiding an infection. This reflects the future disutility of avoiding sickness and death. One problem is that when agents anticipate large infections in the future this value can fall. Jones et al. (2020) call this the fatalism effect.

4 Planner's Problem

We normalize $N = 1$ for simplicity. The planner solves

$$\max U = \sum_{t=0}^{\infty} \beta^t u(C_t, L_t; I_t, D_t)$$

subject to

$$u(C_t, L_t; I_t, D_t) = (1 - D_t) \log(C_t) - (1 - D_t - \kappa I_t) \frac{L_t^{1+\eta}}{1+\eta} - u_\kappa \kappa I_t - u_d D_t$$

and

$$(1 - D_t) C_t = (1 - D_t - \kappa I_t) L_t.$$

The first order conditions for consumption and labor are then (highlighted in red the difference with the decentralized equilibrium)

$$\begin{aligned} C_t : C_t^{-1} &= \lambda_t + 2\lambda_{e,t} e^c C_t, \\ L_t : L_t^\eta &= \lambda_t - 2\lambda_{e,t} e^l L_t. \end{aligned}$$

The marginal utilities of the planner with respect to exposure are twice as high as those of the private sector because of the contagion externalities: private agents only care about how their behavior affect

their own infection risk. They do not care about how their behavior affects the infection risk of others. The envelope condition that differs from that of private agents is

$$V_{I,t} = \kappa \frac{L_t^{1+\eta}}{1+\eta} - \kappa u_\kappa - \kappa \lambda_t L_t + \lambda_{e,t} \kappa e^l L_t^2 - (1-\rho) \lambda_{i,t} - \gamma e_t S_t \lambda_{i,t} - (\delta_t \kappa + \delta'_t \kappa^2 I_t) (\lambda_{d,t} - \lambda_{i,t}).$$

This equation highlights the congestion externality. These externalities determine the planner's incentives to reduce consumption today.

Incentives to Mitigate Let us focus on consumption by setting $e^l = 0$ to understand the incentives to mitigate.

$$C_t^{-1} = \lambda_t + 2\lambda_{e,t} e^c C_t = \lambda_t + 2e^c C_t \gamma I_t S_t (\lambda_{i,t} - \lambda_{s,t})$$

The contemporaneous impact depends on $\gamma I_t S_t$ but the impact is twice as high as in the private case because of the infection externality. As in the decentralized equilibrium, the forward looking effect depends on $\lambda_{i,t} - \lambda_{s,t} = \beta (V_{S,t+1} - V_{I,t+1})$, the future disutility of avoiding sickness and death, which is magnified in the planner solution compared to the decentralized equilibrium because of the potential for congestion in the healthcare system.

5 Calibration

The lack of reliable data to calibrate the contagion model creates a serious challenge and an important limitation. Atkeson (2020a) discusses these difficulties. We calibrate our model at the weekly frequency.

Contagion The SIR block of the model is parameterized as follows. The recovery parameter is set to $\rho = 0.35$. The fraction of infected people who are sick is $\kappa = 0.15$. We normalize $\bar{e} + e^c + e^l = 1$. In our baseline calibration, we set the exposure loading parameters $e^c = e^l = \frac{1}{3}$ which is consistent with the estimation in Ferguson (2020). These parameters imply $e = 1$ at the pre-pandemic levels of consumption and labor (the calibration of production and utility parameters will be described later). The parameter γ is then chosen to target the basic reproduction number (i.e. the average number of people infected by a single infected individual) of $\mathcal{R} = 2$, yielding an estimated value of $\gamma = 0.7$. Finally, to parameterize the fatality rate and the congestion effects, we adopt the following functional form for $\delta(\cdot)$:

$$\delta(\kappa I_t) = \bar{\delta} + \exp(\phi I_t) - 1$$

where the parameter ϕ indexes the strength of the congestion externality. We set $\bar{\delta}$ and ϕ to match two targets for the case fatality rate: a baseline value (i.e. the fraction of infected people who die even in the absence of congestion) of 1% and an ‘extreme’ value (the fraction of people who die $\kappa I = 0.15$ (0.2), i.e. 3% of the population requires medical attention) of 5%. This procedure yields $\bar{\delta} = 0.023$ and $\phi = 3.15$.

Preferences and technology The utility parameter u_d is set to a baseline value of 2. This implies a *flow* disutility from death that is roughly 7 times per capita income. Such large non-monetary costs associated with loss of life are consistent with estimates in the literature and with values used by government entities like the EPA. For example, Greenstone and Nigam (2020) use an estimated value of a statistical life of \$11.5 million (in 2020 dollars) to the household from death. Assuming a rate of return of 5%, this translates into an annual flow value of \$575,000, or roughly 10 times per capita GDP. The flow disutility from sickness u_s is set to equal one-fourth of u_d , i.e. a value of 0.5.

Initial Conditions, Vaccine, and Robustness A time period is interpreted as a week. The discount factor β captures both time discounting and the discovery of a cure/vaccine. We assume for simplicity that a cure and a vaccine arrive randomly together with a constant arrival rate. This is then exactly equivalent to adjusting β . We take a relatively pessimistic case as our baseline, where the combined effect of time discounting and the vaccine is to yield an annual β of 0.8 and a weekly beta of $\beta = (0.8)^{\frac{1}{52}} = 0.9957$.

6 Quantitative Results

Our benchmark exercise uses a large initial infection rate of $i_0 = 1\%$ because it makes the figures easier to read, but this is a large shock. It seems likely that agents and policy makers become aware of the epidemic much earlier so we report simulations starting at $i_0 = 0.1\%$.

Private Response The figures show the results of simulations. We start with the decentralized solution. Figure 1 shows the behavior of the contagion and macro variables in the decentralized equilibrium, under two different assumptions about exposure. The blue line solid shows a situation where infection rates are exogenous, i.e. do not vary with the level of economic activity. Since infection is assumed to be exogenous, agents do not engage in mitigation, i.e., they ignore the pandemic. In fact, labor input rises (the solid line, left panel in the third row in Figure 1), while per-capita consumption

falls by about 2.5% (the dashed line), as able-bodied workers work harder to compensate for the workers who are sick. This is of course not a realistic assumption, but it serves as a useful benchmark for the worst case scenario. In this scenario, eventually about 80% of the population is infected and about 2.5% of the population succumbs to the virus (left panel in the second row in Figure 1). The case mortality rate peaks at 4% roughly 15 weeks after the initial infection because, at the peak, about 15% of the population is infected and the healthcare system is overwhelmed.

The red line describes the case where exposure is endogenous and the household can reduce exposure by cutting back on its consumption and labor supply. As we would expect, this leads to a sharp reduction in economic activity (third row, left panel in Figure 1) by about 10%. Importantly, however, the reduction is gradual, tracking the overall infection rate (it takes almost 17 weeks for consumption and labor to hit their trough). Intuitively, when the fraction of infected people is low (as is the case in the early stages), a reduction in exposure has a small effect on infection risk, relative to the resulting fall in consumption. And since each household does not internalize the effect it has on the future infection rate, it has little incentive to indulge in costly mitigation early on. This dynamic is reflected in the hump-shaped pattern in λ_e (the bottom, left panel in Figure 1). As we will see, this is drastically different in the planner’s problem. The mitigation behavior does lower the cumulative infection and death rates (relative to the exogenous infection risk) down do about 2%.

Optimal Response We now turn to the planner’s solution, depicted in Figures 2. As before, the blue and red lines show the cases of exogenous infection and mitigation. As the red curve in Figure 2 clearly shows, the planner finds it optimal to “flatten the curve” rather dramatically. The peak infection and mortality rates are only slightly higher than their initial levels and well below the decentralized equilibrium levels, as are cumulative fatalities (approximately 0.8%, compared to 2% in the decentralized equilibrium). To achieve this, the planner has to reduce exposure drastically by more than 40% (recall that, in the decentralized equilibrium, exposure bottomed out at 0.8), keeping the basic reproduction number \mathcal{R} from rising much above 1. Of course, this pushes the economy into a deep recession with consumption falling by as much as 40% (third row, left panel in Figure2). More interestingly, the planner chooses to step on the brakes almost immediately, rather wait for infection rates to rise. In fact, the shadow value of exposure (bottom left panel in Figure 2) spikes upon impact and then slowly decays over time, as the number of susceptible people declines.

Early Warning What is the value of an early warning? Suppose agents become aware of the disease at $i_0 = 0.1\%$ instead of 1% as assumed above. This simulation highlights even more the gap between the decentralized outcome and the planner’s solution. The private sector response continues to follow the infection curve. As a result, the outcome barely changes. Private agents do not have the proper incentives to use the early warning.

The planner, on the other hand, continues to front-load her effort and achieves a much better outcome when it receives an early warning. The cumulative fatality rate is only 0.2% instead of 0.8% when it reacts to the disease at a later stage.

7 Conclusions

We propose an extension of the neoclassical model to include contagion dynamics, to study and quantify the tradeoffs of policies that can mitigate the Covid-19 pandemic. Our model reveals two key insights. The first insight is that externalities are massive. The planner acts much more forcefully than private agents. Roughly speaking, under *SIR* dynamics, the planner’s incentives are twice as high as those of private agents. The risk of congestion increases the difference even further. Thus, when private incentives would yield a 10% drop in consumption, the planner engineers an optimal decline of 30% to 40% . One reason why the planner is willing to tolerate such a large decline in consumption is the assumption of complete risk-sharing, i.e. the recession affects all agents equally. Heterogeneity and incomplete markets might make the recession more costly, especially if the burden of the recession falls disproportionately on low-income, low-wealth households, as in Kaplan et al. (2020). Our neoclassical setting also abstracts from demand-related amplifications, as in Guerrieri et al. (2020).

The second key difference is that the planner optimally chooses to front-load her mitigation strategies. As a result, a planner with an early warning does much better than a planner without an early warning. Private agents, on the other hand, waste the value of the early warning because their mitigation efforts are essentially proportional to the current infection rate.

As we write the first draft of this paper there is much uncertainty about the parameters of the disease, and yet decisions must be made. Some of our results speak directly to this dilemma. Atkeson (2020a) points out that, when one does not know the initial number of active cases, it is difficult “to distinguish whether the disease is deadly (1% fatality rate) or milder (0.1% fatality rate).” In our simulations we have considered a deadly disease with a low initial infection rate of $i_0 = 0.1\%$, and a

milder disease with a high initial infection rate of $i_0 = 0.1\%$. Interestingly, in both cases, the planner should implement immediately a strong suppression policy. The main difference is that in the mild case it is optimal to release the lockdown sooner. Assuming that there is enough data 20 weeks after the outbreak to correctly estimate the fatality rate, the planner could implement an optimal response despite the large uncertainty in the key parameter. Jones et al. (2020) study extensions of our baseline setup.

References

Alvarez, Fernando, David Argente, and Francesco Lippi, “A Simple Planning Problem for COVID-19 Lockdown,” March 2020. Working Paper.

Atkeson, Andrew, “How Deadly is COVID-19? Understanding the Difficulties with Estimation,” March 2020. NBER Working Paper No. 26867.

—, “What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios,” March 2020. NBER Working Paper No. 26867.

Baker, Scott R., R.A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis, “How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic,” *University of Chicago, Working Paper*, 2020.

Barro, Robert J, Jose Ursua, and Joanna Weng, “The coronavirus and the Great Influenza epidemic: Lessons from the “Spanish Flu” for the coronavirus’ potential effects on mortality and economic activity,” March 2020. AEI Economics Working Paper 2020-02.

Berger, David, Kyle Herkenhoff, and Simon Mongey, “An SEIR Infectious Disease Model with Testing and Conditional Quarantine,” March 2020. Working Paper.

Correia, Sergio, Stephan Luck, and Emil Verner, “Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu,” March 2020. Working Paper.

Diekmann, Odo and J. Heesterbeek, *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*, John Wiley, 2000.

Eichenbaum, Martin, Sergio Rebelo, and Mathias Trabandt, “The Macroeconomics of Epidemics,” March 2020. Working Paper NWU.

Ferguson, Neil, “Impact of non-pharmaceutical interventions (NPIs) to reduce COVID- 19 mortality and healthcare demand,” March 2020. Imperial College Report.

Greenstone, Michael and Vishan Nigam, “Does Social Distancing Matter?,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, 2020, (2020-26).

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?,” April 2020. NBER Working Paper No. 26918.

Jones, Callum, Thomas Philippon, and Venky Venkateswaran, “Optimal Mitigation Policies in a Pandemic,” March 2020. NYU Working Paper.

Kaplan, Greg, Benjamin Moll, and Gianluca Violante, “Pandemics according to HANK,” April 2020. Working Paper.

Lucas, Robert E. Jr. and Nancy L. Stokey, “Money and Interest in a Cash-in-Advance Economy,” *Econometrica*, 1987, 55 (3), 491–513.

Figure 1: Decentralized Equilibrium

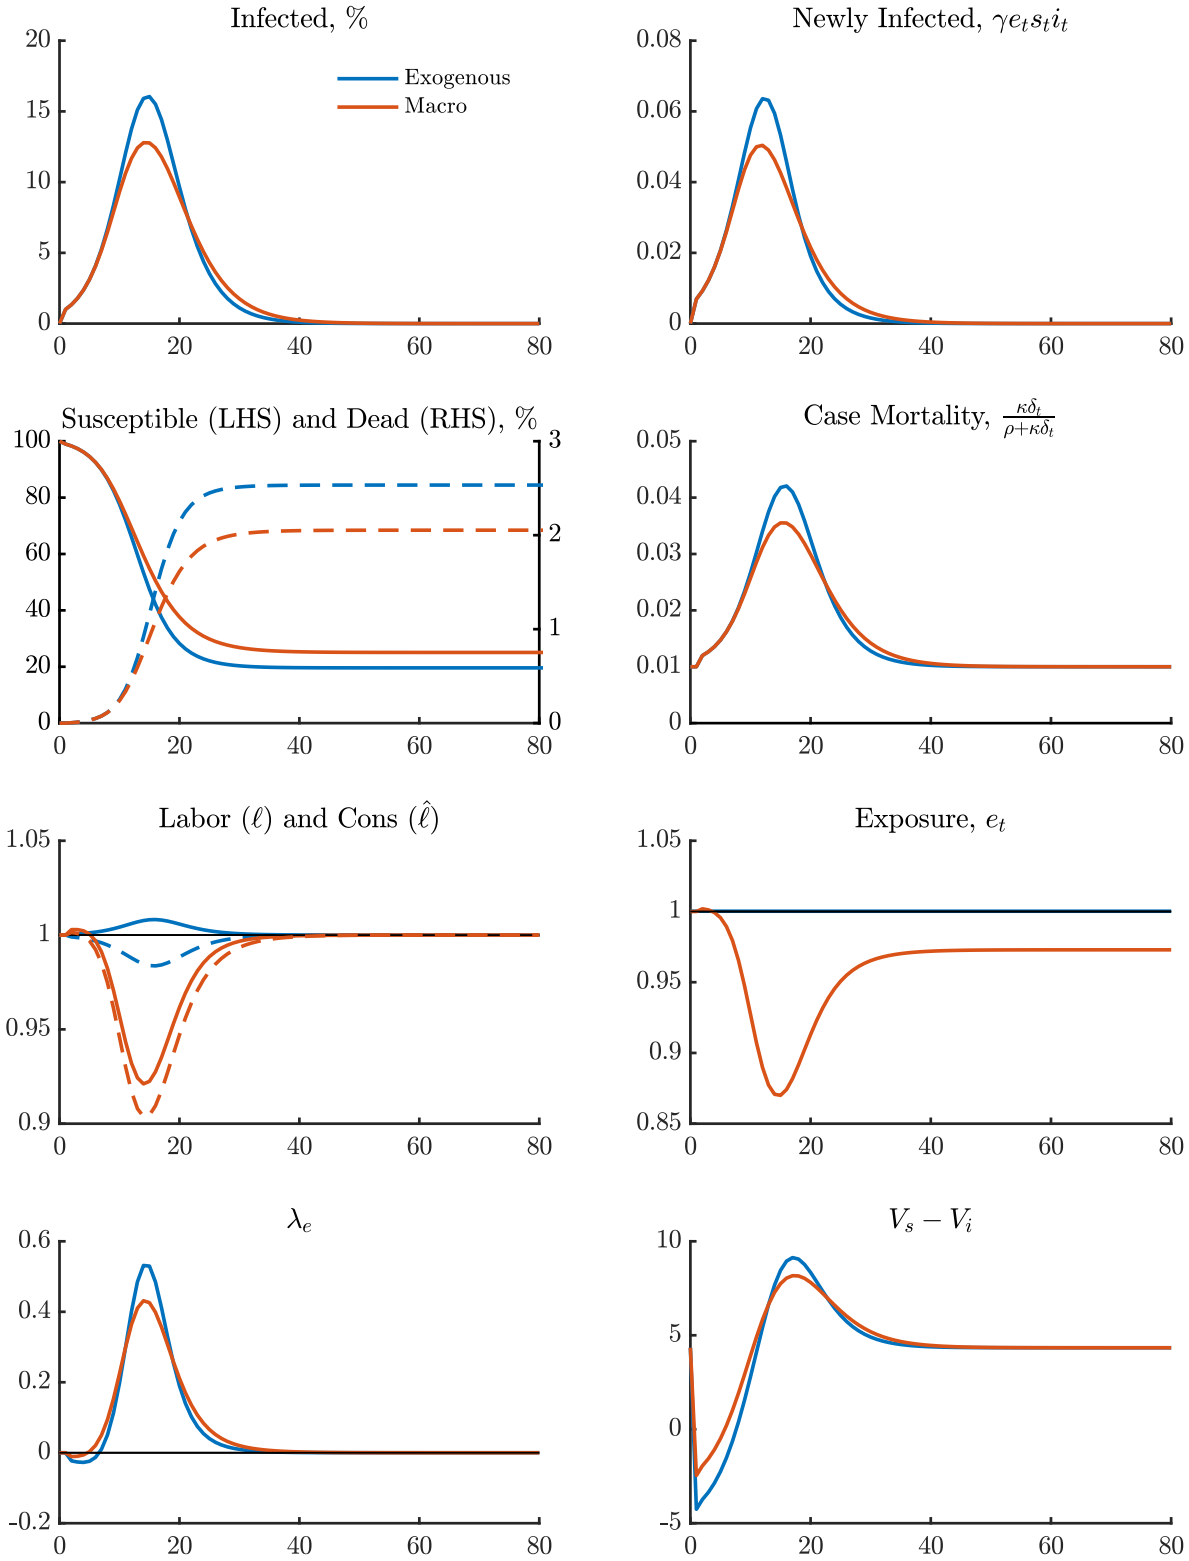


Figure 2: Planner Equilibrium

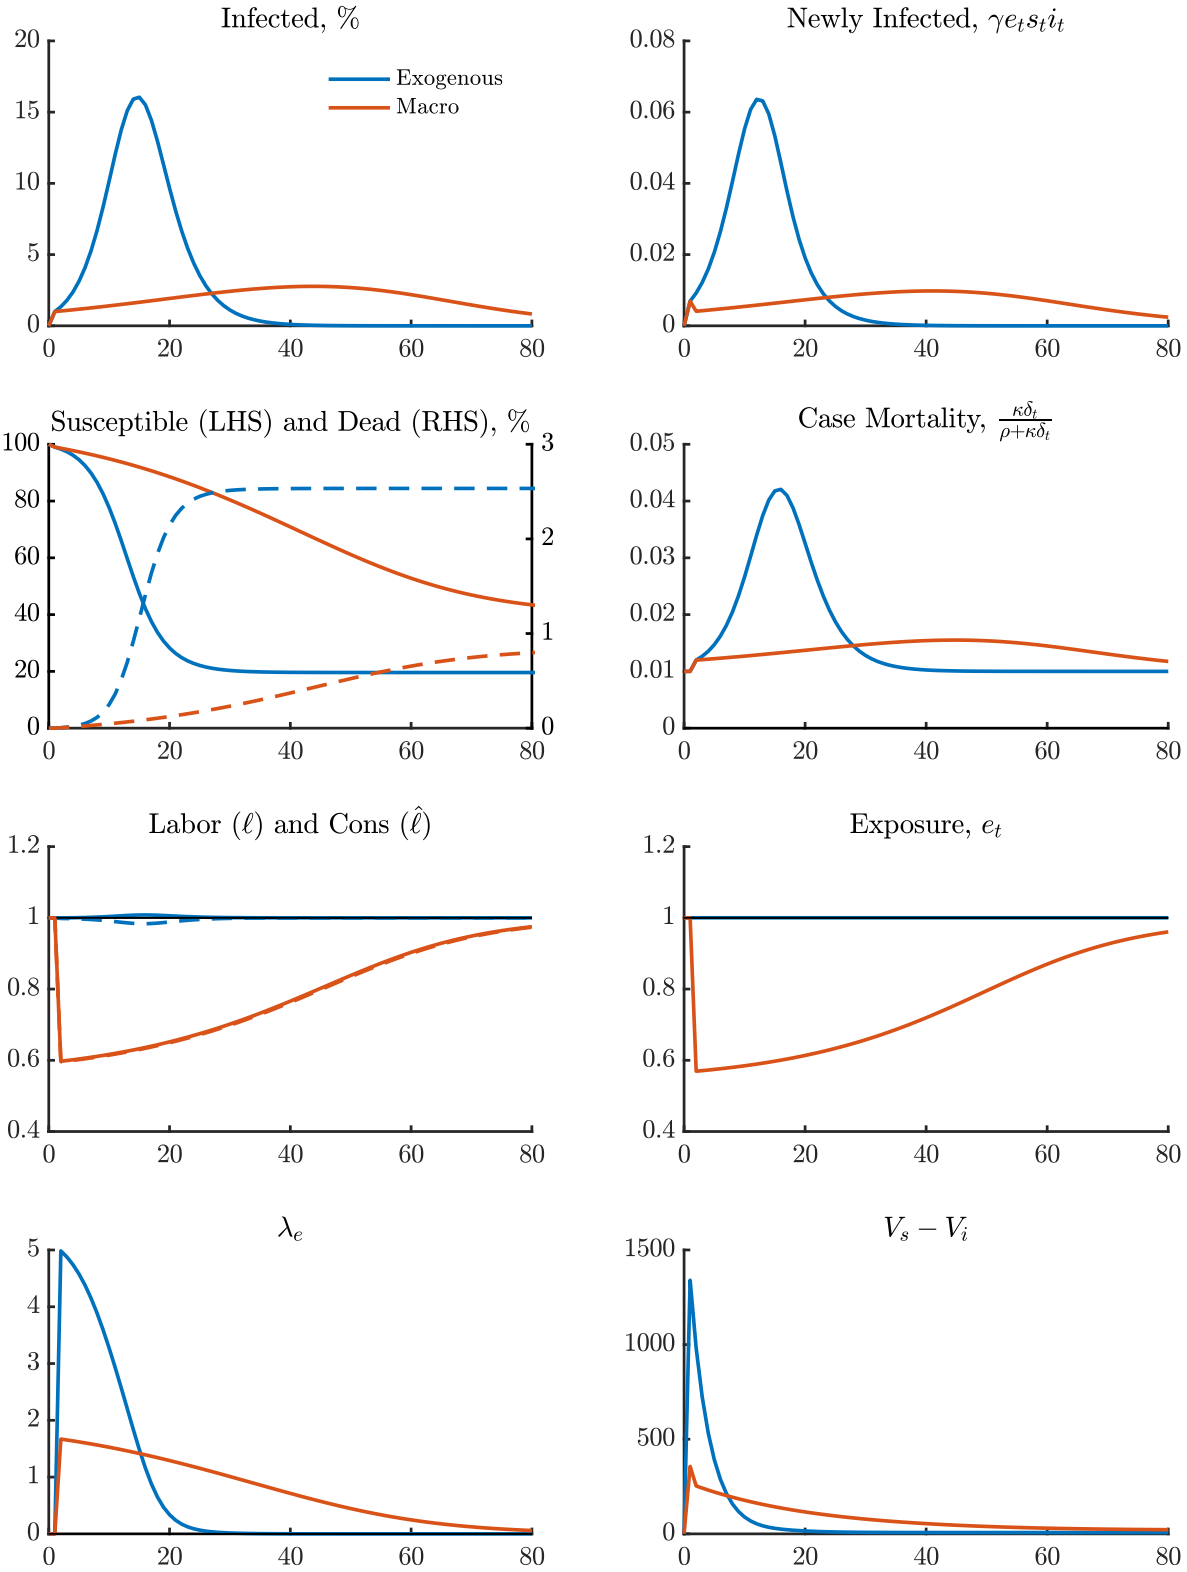


Figure 3: Decentralized Equilibrium, 0.1% Initial Infected

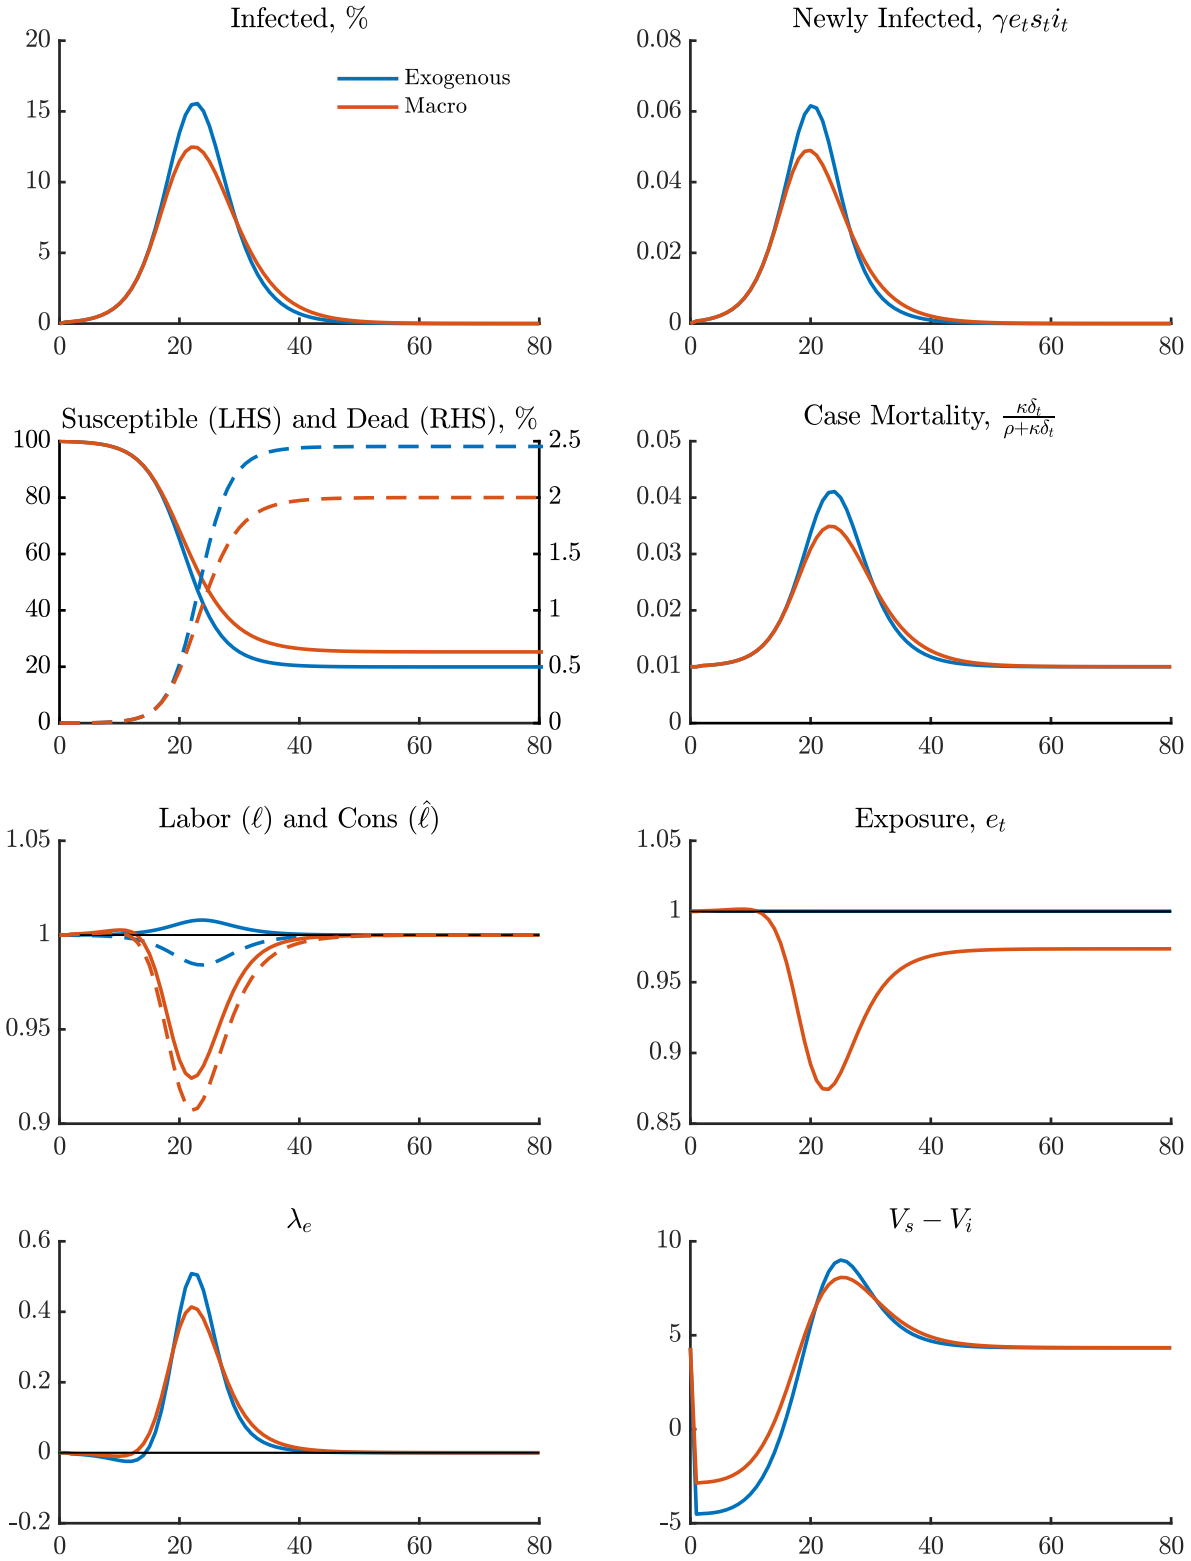


Figure 4: Planner Solution, 0.1% Initial Infected

