NBER WORKING PAPER SERIES

THE MACROECONOMICS OF EPIDEMICS

Martin S. Eichenbaum
Sergio Rebelo
Mathias Trabandt

Working Paper 26882
http://www.nber.org/papers/w26882

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2020

The Macroeconomics of Epidemics
Martin S. Eichenbaum, Sergio Rebelo, and Mathias Trabandt
NBER Working Paper No. 26882
March 2020
JEL No. E1,H0,I1

ABSTRACT

We extend the canonical epidemiology model to study the interaction between economic decisions and epidemics. Our model implies that people's decision to cut back on consumption and work reduces the severity of the epidemic, as measured by total deaths. These decisions exacerbate the size of the recession caused by the epidemic. The competitive equilibrium is not socially optimal because infected people do not fully internalize the effect of their economic decisions on the spread of the virus. In our benchmark scenario, the optimal containment policy increases the severity of the recession but saves roughly 0.6 million lives in the U.S.

Martin S. Eichenbaum
Department of Economics
Northwestern University
2003 Sheridan Road
Evanston, IL 60208
and NBER
eich@northwestern.edu

Sergio Rebelo
Northwestern University
Kellogg School of Management
Department of Finance
Leverone Hall
Evanston, IL 60208-2001
and CEPR
and also NBER
s-rebelo@northwestern.edu

Mathias Trabandt
Freie Universität Berlin
School of Business and Economics
Boltzmannstrasse 20
14195 Berlin
Germany
and DIW and IWH
mathias.trabandt@gmail.com

1 Introduction

As the COVID-19 virus spreads throughout the world, governments are struggling with how to understand and manage the epidemic. Epidemiology models have been widely used to predict the course of the epidemic. While these models are very useful, they do have an important shortcoming: they do not allow for the interaction between economic decisions and rates of infection.

Policy makers certainly appreciate this interaction. For example, in their March 19, 2020 Financial Times op ed Ben Bernanke and Janet Yellen write that

> "In the near term, public health objectives necessitate people staying home from shopping and work, especially if they are sick or at risk. So production and spending must inevitably decline for a time."

In this paper, we extend the classic SIR model proposed by Kermack and McKendrick (1927) to study the interaction between economic decisions and epidemic dynamics.[1] Our model makes clear that people's decisions to cut back on consumption and work reduce the severity of the epidemic as measured by total deaths. These same decisions exacerbate the size of the recession caused by the epidemic.

In our model, an epidemic has both aggregate demand and aggregate supply effects. The supply effect arises because the epidemic exposes people who are working to the virus. People react to that risk by reducing their labor supply. The demand effect arises because the epidemic exposes people who are purchasing consumption goods to the virus. People react to that risk by reducing their consumption. The supply and demand effects work together to generate a large, persistent recession.

The competitive equilibrium is not Pareto optimal because people infected with the virus do not fully internalize the effect of their consumption and work decisions on the spread of the virus. To be clear, this market failure does not reflect a lack of good intentions or irrationality on the part of infected people. It simply reflects the fact that each individual agent takes economy-wide infection rates as given.

A natural question is: what policies should the government pursue to deal with the infection externality? We focus on containment policies that reduce consumption and hours worked. By reducing economic interactions among agents, these policies exacerbate the

[1] SIR is an acronym for susceptible, infected, and recovered or removed.

recession but raise welfare by reducing the death toll caused by the epidemic. We find that it is optimal to introduce large-scale containment measures that result in a sharp, sustained drop in aggregate output. This optimal containment policy saves about $600,000$ lives in the U.S.

To make the intuition for our results as transparent as possible, we use a relatively simple model. A cost of that simplicity is that we cannot study many important, epidemic-related policy issues. For example, we do not consider polices aimed at mitigating the economic hardships suffered by households and businesses. Such policies include fiscal transfers to people and loans to keep firms from going bankrupt. We also do not study policies aimed at maintaining the liquidity and health of financial markets.

Finally, we abstract from nominal rigidities which could play an important role in determining the short-run response of the economy to an epidemic. For example, if prices are sticky, a given fall in the demand for consumption would generate a larger recession. Other things equal, a larger recession would mitigate the spread of the infection. We plan to address these important issues in future work. But we are confident that the central message from our current analysis will be robust: there is an inevitable trade-off between the severity of the recession and the health consequences of the epidemic.

Our point of departure is the canonical SIR model proposed by Kermack and McKendrick (1927). In this model the transition probability between health states are exogenous parameters. We modify the model by assuming that purchasing consumption goods and working brings people in contact with each other. These activities raise the probability that the infection spreads. We refer to the resulting framework as the SIR-macro model.

We choose parameters so that the Kermack-McKendrick SIR model is consistent with the scenario outlined by Angela Merkel in her March 11, 2020 speech.[2] According to this scenario, "60 to 70 percent of the population will be infected as long as this remains the situation." The SIR model implies that the share of the initial population infected peaks at 8.4 percent. Applying this scenario to the U.S. implies that roughly 215 million Americans eventually become infected and 2.2 million people die. When we embed the SIR model in a simple general equilibrium framework, we find that the epidemic causes a relatively mild recession. Aggregate consumption falls by roughly 2 percent from peak to trough, with the latter occurring 29 weeks after the onset of the infection. In the long-run, population and real GDP decline permanently by 0.65 percent reflecting the death toll from the epidemic.

[2] "Merkel Gives Germans a Hard Truth About the Corona Virus," New York Times, March 11, 2020.

The impact of economic activity on transition probabilities in the SIR-macro model, substantially changes the dynamics of the epidemic and its economic impact. Relative to the SIR model, the SIR-macro model implies a sharper recession and fewer deaths. The peak to trough decline in aggregate consumption is more than four times as large as in the standard SIR model (9.3 versus 2 percent). This larger decline in economic activity reduces the infection peak (5.1 percent versus 8.4 percent) as well as the percentage of the population that becomes infected (52.7 versus 65 percent). Critically, the total number of U.S. deaths caused by the epidemic falls from 2.2 to 1.74 million.

How do epidemics end? In both the SIR and SIR-macro models, epidemics end when a sufficiently high fraction of the population acquires immunity. Absent treatments or vaccines, the only way to acquire immunity is to become infected and recover. Sadly, this process involves the death of many people who never recover from an infection. This property of epidemics serves as an important backdrop for our discussion of optimal policy.

Given the negative externalities from consumption and work, the optimal policy in the SIR-macro model is to curtail economic activity. In all versions of our model, it is optimal for policymakers to avoid recurrent epidemics. So, absent a vaccine or treatment, they must allow a sufficiently high fraction of the population to be become infected and recover. The key question is: what is the optimal way to reach that fraction?

In the SIR-macro model it is possible to prevent the infection from spreading by adopting large, permanent containment measures. The problem with this approach is that the population never reaches the critical level of immunity to avoid a recurrence of the epidemic. Under these circumstances, infections would recur as soon as containment is relaxed. The optimal policy in this world is to build up the fraction of the population that is immune, curtailing consumption when externalities are large, that is when the number of infected people is high. Such a policy involves gradually ramping up containment measures as infections rise and slowly relaxing them as new infections wane and the population approaches the critical immunity level.

An important concern in many countries is that the healthcare system will be overwhelmed by a large number of infected people. To analyze this scenario, we consider a version of our model in which the mortality rate is an increasing function of the number of people infected. We find that the competitive equilibrium involves a much larger recession, as people internalize the higher mortality rates. People cut back more aggressively on consumption and work to reduce the probability of being infected. As a result, fewer people are

infected in the competitive equilibrium but more people die. The optimal policy involves a much more aggressive response than in the baseline SIR-macro economy. The reason is that the cost of the externality is much larger since a larger fraction of the infected population dies.

How does the possibility of an effective treatment being discovered change our results? The qualitative implications are clear: people become more willing to engage in market activities because the expected cost of being infected is smaller. So, along a path in which treatment is not actually discovered, the recession induced by the epidemic is less severe. Sadly, along such a path, the total number of infected people and the death toll rise relative to the baseline SIR-macro model. That said, the quantitative difference of this model and the baseline SIR-macro model is quite small, both with respect to the competitive equilibrium and the optimal containment policy.

How does the possibility of a vaccine being discovered change our results? Vaccines don't cure infected people but they do prevent susceptible people from becoming infected. In contrast, treatments cure infected people but do not prevent future infections. These differences are not very important for the competitive equilibrium. But they have very different implications for optimal policy. With vaccines as a possibility, it is optimal to *immediately* introduce severe containment measures to minimize deaths. Those containment measures cause a large recession. But this recession is worth incurring in the hope that the vaccination arrives before many people get infected.

Our paper is organized as follows. In section 2, we describe both the SIR and the SIR-macro model. In section 3, we describe the versions of the model that consider medical preparedness and the possibility of effective treatment and vaccines being discovered. In section 4, we discuss the properties of the competitive equilibrium in different variants of our model. In section 5, we solve the Ramsey policy problems and analyze their implications for the containment of the spread of the virus and for economic activity. Section 6 concludes.

2 The SIR-macro model

In this section, we describe the economy before the start of the epidemic. We then present the SIR-macro model.

2.1 The pre-infection economy

The economy is populated by a continuum of identical agents with measure one. Prior to the start of the epidemic, all agents are identical and maximize the objective function:

$$U = \sum_{t=0}^{\infty} \beta^t u(c_t, n_t).$$

Here $\beta \in (0,1)$ denotes the discount factor and c_t and n_t denote consumption and hours worked, respectively. For simplicity, we assume that momentary utility takes the form

$$u(c_t, n_t) = \ln c_t - \frac{\theta}{2} n_t^2.$$

The budget constraint of the representative agent is:

$$(1 + \mu_{ct}) c_t = w_t n_t + \Gamma_t.$$

Here, w_t denotes the real wage rate, μ_{ct} is the tax rate on consumption, and Γ_t denotes lump-sum transfers from the government. As discussed below, we think of μ_{ct} as a proxy for containment measures aimed at reducing social interactions. For this reason, we refer to μ_{ct} as the containment rate. The first-order condition for the representative-agent's problem is:

$$(1 + \mu_{ct}) \theta n_t = c_t^{-1} w_t.$$

There is a continuum of competitive representative firms of unit measure that produce consumption goods (C_t) using hours worked (N_t) according to the technology:

$$C_t = A N_t.$$

The firm chooses hours worked to maximize its time-t profits Π_t:

$$\Pi_t = A N_t - w_t N_t.$$

The government's budget constraint is given by

$$\mu_{ct} c_t = \Gamma_t.$$

In equilibrium, $n_t = N_t$ and $c_t = C_t$.

2.2 The outbreak of an epidemic

Epidemiology models generally assume that the probabilities governing the transition between different states of health are exogenous with respect to economic decisions. We modify the classic SIR model proposed by Kermack and McKendrick (1927) so that these transition probabilities depend on people's economic decisions. Since purchasing consumption goods or working brings people into contact with each other, we assume that the probability of becoming infected depends on these activities.

The population is divided into four groups: susceptible (people who have not yet been exposed to the disease), infected (people who contracted the disease), recovered (people who survived the disease and acquired immunity), and deceased (people who died from the disease). The fractions of people in these four groups are denoted by S_t, I_t, R_t and D_t, respectively. The number of newly infected people is denoted by T_t.

Susceptible people can become infected in three ways. First, they can meet infected people while purchasing consumption goods. The number of newly infected people that results from these interactions is given by $\pi_{s1}(S_t C_t^S)\left(I_t C_t^I\right)$. The terms $S_t C_t^S$ and $I_t C_t^I$ represent total consumption expenditures by susceptible and infected people, respectively. The parameter π_{s1} reflects both the amount of time spent shopping and the probability of becoming infected as a result of that activity. In reality, different types of consumption involve different amounts of contact with people. For example, attending a rock concert is much more contact intensive than going to a grocery store. For simplicity we abstract from this type of heterogeneity.

Second, susceptible and infected people can meet at work. The number of newly infected people that results from interactions at work is given by $\pi_{s2}(S_t N_t^S)\left(I_t N_t^I\right)$. The terms $S_t N_t^S$ and $I_t N_t^I$ represent total hours worked by susceptible and infected people, respectively. The parameter π_{s2} reflects the probability of becoming infected as a result of work interactions. We recognize that different jobs require different amounts of contact with people. For example, working as a dentist or a waiter is much more contact intensive than writing software. Again, for simplicity, we abstract from this source of heterogeneity.

Third, susceptible and infected people can meet in ways not directly related to consuming or working, for example meeting a neighbor or touching a contaminated surface. The number of random meetings between infected and susceptible people is $S_t I_t$. These meetings result in $\pi_{s3} S_t I_t$ newly infected people.

The total number of newly infected people is given by:

$$T_t = \pi_{s1}(S_t C_t^S)\left(I_t C_t^I\right) + \pi_{s2}(S_t N_t^S)\left(I_t N_t^I\right) + \pi_{s3} S_t I_t. \tag{1}$$

Kermack and McKendrick's (1927) SIR model is a special case of our model in which the propagation of the disease is unrelated to economic activity ($\pi_{s1} = 0$, $\pi_{s2} = 0$).

The number of susceptible people at time $t+1$ is equal to the number of susceptible people at time t minus the number of susceptible people that got infected at time t:

$$S_{t+1} = S_t - T_t. \tag{2}$$

The number of infected people at time $t+1$ is equal to the number of infected people at time t plus the number of newly infected (T_t) minus the number infected people that recovered ($\pi_r I_t$) and the number of infected people that died ($\pi_d I_t$):

$$I_{t+1} = I_t + T_t - (\pi_r + \pi_d) I_t. \tag{3}$$

Here, π_r is the rate at which infected people recover from the infection and π_d is the mortality rate, that is the probability that an infected person dies.

The number of recovered people at time $t+1$ is the number of recovered people at time t plus the number of infected people who just recovered ($\pi_r I_t$):

$$R_{t+1} = R_t + \pi_r I_t. \tag{4}$$

Finally, the number of deceased people at time $t+1$ is the number of deceased people at time t plus the number of new deaths ($\pi_d I_t$):

$$D_{t+1} = D_t + \pi_d I_t. \tag{5}$$

Total population, Pop_{t+1}, evolves according to:

$$Pop_{t+1} = Pop_t - D_t,$$

with $Pop_0 = 1$.

We assume that at time zero a fraction ε of susceptible people is infected by a virus through zoonotic exposure, that is the virus is directly transmitted from animals to humans,

$$I_0 = \varepsilon,$$

$$S_0 = 1 - \varepsilon.$$

All agents are aware of the initial infection and understand the laws of motion governing population health dynamics.

We now describe the optimization problem of different types of people in the economy. The variable U_t^j denotes the time-t lifetime utility of a type-j agent ($j = s, i, r$). The budget constraint of a type-j person is

$$(1 + \mu_{ct})c_t^j = w_t \phi^j n_t^j + \Gamma_t, \tag{6}$$

where c_t^j and n_t^j denote the consumption and hours worker of agent j, respectively. The parameter governing labor productivity, ϕ^j, is equal to one for susceptible and recovered people ($\phi^s = \phi^r = 1$) and less than one for infected people ($\phi^i < 1$).

The budget constraint (6) embodies the assumption that there is no way for agents to pool risk associated with the infection. Going to the opposite extreme and assuming complete markets considerably complicates the analysis without necessarily making the model more realistic.

Susceptible people The lifetime utility of a susceptible person, U_t^s, is

$$U_t^s = u(c_t^s, n_t^s) + \beta \left[(1 - \tau_t) U_{t+1}^s + \tau_t U_{t+1}^i \right]. \tag{7}$$

Here, the variable τ_t represents the probability that a susceptible person becomes infected:

$$\tau_t = \pi_{s1} c_t^s \left(I_t C_t^I \right) + \pi_{s2} n_t^s \left(I_t N_t^I \right) + \pi_{s3} I_t. \tag{8}$$

Susceptible people take as given aggregate variables like $I_t C_t^I$ and $I_t N_t^I$. Critically, they understand that consuming and working less reduces the probability of becoming infected.

The first-order conditions for consumption and hours worked are:

$$u_1(c_t^s, n_t^s) - (1 + \mu_{ct})\lambda_{bt}^s + \lambda_{\tau t} \pi_{s1} \left(I_t C_t^I \right) = 0,$$

$$u_2(c_t^s, n_t^s) + A\lambda_{bt}^s + \lambda_{\tau t} \pi_{s2} \left(I_t N_t^I \right) = 0.$$

Here λ_{bt}^s and $\lambda_{\tau t}$ are the Lagrange multipliers associated with the constraints (6) and (8), respectively.

The first-order condition for τ_t is:

$$\beta \left(U_{t+1}^i - U_{t+1}^s \right) - \lambda_{\tau t} = 0. \tag{9}$$

Infected people The lifetime utility of an infected person, U_t^i, is

$$U_t^i = u(c_t^i, n_t^i) + \beta \left[(1 - \pi_r - \pi_d) U_{t+1}^i + \pi_r U_{t+1}^r + \pi_d \times 0 \right]. \tag{10}$$

The first-order conditions for consumption and hours worked are given by

$$u_1(c_t^i, n_t^i) = \lambda_{bt}^i (1 + \mu_{ct}),$$

$$u_2(c_t^i, n_t^i) = -\phi^i A \lambda_{bt}^i,$$

where λ_{bt}^i is the Lagrange multiplier associated with the constraint (6).

Recovered people The lifetime utility of a recovered person, U_t^r, is

$$U_t^r = u(c_t^r, n_t^r) + \beta U_{t+1}^r. \tag{11}$$

The first-order conditions for consumption and hours worked are:

$$u_1(c_t^r, n_t^r) = \lambda_{bt}^r (1 + \mu_{ct})$$

$$u_2(c_t^r, n_t^r) = -A \lambda_{bt}^r$$

where λ_{bt}^r is the Lagrange multiplier associated with the constraint (6).

Equilibrium In equilibrium, each person solves their maximization problem and the government budget constraint is satisfied. In addition, the goods and labor markets clear.

$$S_t C_t^s + I_t C_t^i + R_t C_t^r = C_t,$$

$$S_t N_t^s + I_t N_t^i + R_t N_t^r = N_t.$$

In the appendix we describe our algorithm for computing the equilibrium.

3 Medical preparedness, treatments and vaccines

In this section we extend the SIR-macro model in three ways. First, we allow for the possibility that the mortality rate increases as the number of infections rises. Second, we allow for the probabilistic development of a cure for the disease. Third, we allow for the probabilistic development of a vaccine that inoculates susceptible people against the virus.

3.1 The medical preparedness model

In our benchmark model we abstracted from the possibility that the efficacy of the healthcare system will deteriorate if a substantial fraction of the population becomes infected. A simple way to model this scenario is to assume that the mortality rate depends on the number of infected people, I_t:

$$\pi_{dt} = \pi_d + \kappa I_t^2.$$

This functional form implies that the mortality rate is a convex function of the fraction of the population that becomes infected. The benchmark SIR-macro corresponds to the special case of $\kappa = 0$.

3.2 The treatment model

The benchmark SIR-macro model abstracts from the possibility that an effective treatment against the virus will be developed. Suppose instead that an effective treatment is discovered with probability δ each period. Once discovered, treatment is provided to all infected people in the period of discovery and all subsequent periods transforming them into recovered people. As a result, the number of new deaths from the disease goes to zero.

The lifetime utility of an infected person before the treatment becomes available is:

$$U_t^i = u(c_t^i, n_t^i) + (1-\delta)\left[(1-\pi_r-\pi_d)\,\beta U_{t+1}^i + \pi_r \beta U_{t+1}^r + \pi_d \times 0\right] + \beta\delta U_{t+1}^r.$$

This expression reflects the fact that with probability $1-\delta$ a person who is infected at time t remains so at time $t+1$. With probability δ this person receives treatment and becomes recovered.

We now discuss the impact of an effective treatment on population dynamics. Before the treatment is discovered, population dynamics evolve according to equations (1), (2), (3), (4), and (5). Suppose that the treatment is discovered at the beginning of time t^*. Then, all infected people become recovered. The number of deceased stabilizes once the treatment arrives:

$$D_t = D_{t^*} \text{ for } t \geq t^*.$$

Since anyone can be instantly cured, we normalize the number of susceptible and infected people to zero for $t > t^*$. The number of recovered people is given by

$$R_t = 1 - D_t.$$

3.3 The vaccination model

The benchmark SIR-macro model abstracts from the possibility that a vaccine against the virus will be developed. Suppose instead that a vaccine is discovered with probability δ per period. Once discovered, the vaccine is provided to all susceptible people in the period of discovery and in all subsequent periods.

The lifetime utility of a susceptible person is given by

$$U_t^s = u(c_t^s, n_t^s) + (1-\delta)\left[(1-\tau_t)\,\beta U_{t+1}^s + \tau_t \beta U_{t+1}^i\right] + \delta\beta U_{t+1}^r. \tag{12}$$

This expression reflects the fact that with probability $1-\delta$ a person who is susceptible at time t remains so at time $t+1$. With probability δ this person is vaccinated and becomes immune to the disease. So, at time $t+1$, this person's health situation is identical to that of a recovered person. The vaccine has no impact on people who were infected or have recovered. The lifetime utilities of infected and recovered people person are given by (10) and (11), respectively.

We now discuss the impact of vaccinations on population dynamics. Before the vaccine is discovered, these dynamics evolve according to equations (1), (2), (3), (4), and (5). Suppose that the vaccine is discovered at the beginning of time t^*. Then, all susceptible people become recovered. Since no one is susceptible, there are no new infections.

Denote the number of susceptible and recovered people right after a vaccine is introduced at time t^* by S'_{t^*} and R'_{t^*}. The value of these variables are

$$S'_{t^*} = 0$$

$$R'_{t^*} = R_{t^*} + S_{t^*}.$$

For $t \geq t^*$ we have

$$R_{t+1} = \begin{cases} R'_t + \pi_r I_t & \text{for } t = t^* \\ R_t + \pi_r I_t & \text{for } t > t^*. \end{cases}$$

The laws of motion for I_t and D_t are given by (3) and (5).

4 Competitive equilibrium

In this section, we discuss the properties of the competitive equilibrium via a series of numerical exercises. In the first subsection, we describe our parameter values. In the second

and third subsections, we discuss how the economy responds to an epidemic in the SIR and SIR-macro models, respectively. In the fourth subsection, we discuss the implications of medical preparedness. Finally, in the fifth subsection, we discuss the effects of treatments and vaccines.

4.1 Parameter values

Each time period corresponds to a week. While acknowledging considerable uncertainty about infection, recovery and mortality rates, we choose π_d so that the mortality rate is one percent. This value is equal to the estimate for the mortality rate from COVID-19 reported by the World Health Organization on March 16, 2020 for South Korea.[3] Estimates of the mortality rate based on South Korean data are relatively reliable because that country has the world's highest per capita test rates for COVID-19 (Pueyo (2020)). Estimates of mortality rates based on data from other countries are probably biased upwards because the number of infected people is likely to be underestimated. As in Atkeson (2020), we assume that it takes on average 18 days to either recover or die from the infection. Since our model is weekly, we set $\pi_r + \pi_d = 7/18$. A one percent mortality rate for infected people implies $\pi_d = 7 \times 0.01/18$.

We use the standard SIR model to choose values for π_{s1}, π_{s2}, and π_{s3}. Our procedure is as follows. First, we assume each of the three terms in (1), representing the different ways in which people can get infected, accounts for 1/3 of the value of T_0. This assumption yields two independent restrictions on the parameters:

$$\frac{\pi_{s1}C^2}{\pi_{s1}C^2 + \pi_{s2}N^2 + \pi_{s3}} = \frac{\pi_{s2}N^2}{\pi_{s1}C^2 + \pi_{s2}N^2 + \pi_{s3}} = \frac{1}{3}.$$

Here, C and N are consumption and hours worked in the pre-infection steady state. In addition, we assume that in the limit of 65 percent of the population either recovers from the infection or dies. This assumption corresponds to the Merkel scenario discussed in the introduction. With this scenario we have a third restriction that allows us to compute π_{s1}, π_{s2}, and π_{s3}. The resulting values for these parameters are 0.0185, 1.8496, and 0.2055, respectively.

The initial population is normalized to one. The number of people that are initially infected, ε, is 0.001. We assume that $A = 10$, a value that ensures that U_i is roughly one percent lower than U_s.

[3] This estimate is roughly 8 times larger than the average flu death rate in the U.S.

The parameter θ, which controls the disutility of labor, is set to 9 so that the number of hours worked in the pre-epidemic steady state is 1/3. We set ϕ^i, the parameter that controls the relative productivity of infected people is 0.8. This value is consistent with the notion that symptomatic agents don't work and the assumption that 80 percent of infected agents are asymptomatic (Pueyo (2020)). We set the discount factor β to be consistent with a 2 percent annual interest rate, $\beta = 0.98^{1/52}$. In the baseline SIR-macro model μ_{ct} is equal to zero.

In the medical preparedness model, we fix κ to 13, which implies a peak mortality rate of 4 percent. This value coincides with the current WHO estimates of the mortality rate in China, a country where the number of infections seems to be close to its peak.

In both the treatment and vaccination models we set $\delta = 1/52$ which implies that it takes on average 52 weeks for these medical discoveries to become available.

4.1.1 The model's basic reproduction number

A widely used statistic used to diagnose the severity of an epidemic is the "basic reproduction number," $\mathcal{R}_0$. This statistic is the total number of infections caused by one infected person (with measure zero) in his or her lifetime in a population where everybody is susceptible ($S_t = 1$). The higher is the value of $\mathcal{R}_0$, the faster is the spread of the virus.

The average rate of infection, which we denote by γ, in our model is the ratio of the number of newly infected people to the total number of infected people. The value of γ is equal to T_0/I_0. The expected number of infections caused by a single infected person is

$$\gamma + (1 - \pi_r - \pi_d)\gamma + (1 - \pi_r - \pi_d)^2\gamma + ... = \frac{\gamma}{\pi_r + \pi_d}.$$

In this expression, $(1 - \pi_r - \pi_d)^t$ is the probability that the infected person reaches period t without recovering or dying.

The value of $\mathcal{R}_0$ in the SIR and benchmark SIR-macro models is 1.58 and 1.48, respectively. These values are lower than current point estimates of $\mathcal{R}_0$ for COVID-19, but consistent with the evidence taking sampling uncertainty into account. For example, Riou and Althaus (2020) report a point estimate of 2.2 with a 90 percent confidence interval of 1.4 to 3.8.

4.2 The SIR model

The black dashed lines in Figure 1 display the equilibrium population dynamics implied by the SIR model. Note that the share of the initial population that is infected peaks at 8.4 percent in week 28. Thereafter, this share falls because there are less susceptible people to infect. Eventually, 65 percent of the population becomes infected. Assuming a U.S. population of roughly 330 million people, this scenario implies that roughly 215 million Americans eventually become infected. A mortality rate of one percent ($\pi_d = 0.01$) implies that roughly 2.15 million people die.

Figure 1 shows that the epidemic induces a recession: aggregate consumption falls by roughly 2 percent from peak to trough. This fall reflects two factors. First and foremost, the virus causes infected people to be less productive at work ($\phi^i = 0.8$). The associated negative income effect lowers the consumption of those who are infected. The dynamic behavior of aggregate consumption mimics the share of infected agents in the overall population. Second, the death toll caused by the epidemic permanently reduces the size of the work force.

Since production is constant returns to scale, per capita income is the same in the post- and pre-epidemic steady states. In the post-epidemic steady state, population and real GDP are both 0.65 percent lower than in the initial steady state.

4.3 The SIR-macro model

In the SIR model economic decisions about consumption and work don't influence the dynamics of the epidemic. In the SIR-macro model, susceptible households can lower the probability of being infected by reducing their consumption and hours worked. The solid blue lines in Figure 1 show how the epidemic unfolds in this model.

Note that the share of the initial population that is infected peaks at 5.1 percent in period 32. The peak is substantially smaller and occurs somewhat later than the corresponding peak in the SIR model. Eventually, 52.7 percent of the population becomes infected. So, for the U.S., roughly 174 million people eventually become infected and 1.74 million people die.

Figure 1 shows that the infection is less severe in the SIR-macro model than in the SIR model. The reason is that in the SIR-macro model susceptible people severely reduce their consumption and hours worked to lower the probability of being infected. Consistent with Figure 2, there are no offsetting effects arising from the behavior of recovered and infected people because they behave as in the SIR model.

Consistent with these observations, the recession is much more severe in the SIR-macro model: the peak to trough fall in aggregate consumption is 9.3 percent, more than 4.5 larger than in the SIR model.

For similar reasons, the dynamics and magnitude of the drop in hours work is very different in the two models. In the SIR model, hours worked decline smoothly falling by 0.65 percent in the post-epidemic steady state. This decline entirely reflects the impact of the death toll on the workforce.

In the SIR-macro model, hours worked follow a U-shaped pattern. The peak decline of 8.25 percent occurs in period 32. Thereafter, aggregate hours rise converging to a new steady state from below. These dynamics are driven by the labor-supply decisions of susceptible agents. Interestingly, the long-run decline in hours worked is lower in the SIR-macro model (0.53 percent) than in the SIR model (0.65 percent). The reason is that fewer people die in the epidemic so the population falls by less in the SIR-macro model than in the SIR model.

Figure 3 shows the competitive equilibrium and the optimal containment policy in the SIR-macro model. We return to this Figure in the next section.

4.4 Medical preparedness model

The red dashed-dotted lines in Figure 4 show that the competitive equilibrium with an endogenous mortality rate involves a much larger recession than in the baseline SIR-macro model (blue solid lines). The reason is that people internalize the higher mortality rates associated with an healthcare system that is overburdened with infected people. Since the costs of becoming infected are much higher, people cut back on consumption and work to reduce the probability of becoming infected. The net result is that fewer people are infected but more people die.

4.5 The treatment and vaccines models

As discussed in the introduction, the possibility of treatment and vaccination have similar qualitative effects on the competitive equilibrium. Compared to the baseline SIR-macro model people become more willing to engage in market activities. The reason is that the expected costs associated with being infected are smaller. Because of this change in behavior, the recession is less severe. In Figures 5 and 6 the blue-solid and red-dashed-dotted lines virtually coincide. So, in practice the quantitative effect of the possibility of treatments or vaccinations on the competitive equilibrium is quite small.

5 Economic policy

The competitive equilibrium of our model economy is not Pareto optimal. There is a classic externality associated with the behavior of infected agents. Because agents are atomistic, they don't take into account the impact of their actions on the infection and death rates of other agents. In this section, we consider a simple Ramsey problem designed to deal with this externality.

As with any Ramsey problem, we must take a stand on the policy instruments available. In reality, there are many ways in which governments can reduce social interactions. Examples of containment measures include shelter-in-place laws and shutting down of restaurants and bars. Analogous to Farhi and Werning's (2012) treatment of capital controls, we model these measures as a tax on consumption, the proceeds of which are rebated lump sum to all agents. We refer to this tax as the containment rate.

We compute the optimal sequence of 150 containment rates $\{\mu_{ct}\}_{t=0}^{149}$ that maximize social welfare,U_0, defined as a weighted average of the lifetime utility of the different agents:

$$U_0 = s_0 U_0^s + i_0 U_0^i + r_0 U_0^r.$$

Given the sequence of containment rates, we solve for the competitive equilibrium and evaluate the social welfare function. We iterate on this sequence until we find the optimum.

Figure 3 displays our results. First, it is optimal to escalate containment measures gradually over time. The optimal containment rate rises from 2.35 percent in period zero to a peak value of 47.3 percent in period 40. The rise in containment rates roughly parallels the dynamics of the infection rate itself. The basic intuition is as follows. Containment measures internalize the externality caused by the behavior of infected people. So, as the number of infected people rises it is optimal to intensify containment measures. For example, at time zero very few people are infected, so the externality is relatively unimportant. A high containment rate at time zero would have a high social cost relative to the benefit. As the infection rises, the externality becomes important and the optimal containment rate rises.

The optimal containment policy greatly reduces the peak level of infections from 5.1 to 2.5 percent reducing the death toll from 0.53 to 0.36 percent of the initial population. For a country like the U.S., this reduction represents roughly 0.6 million lives saved. This beneficial outcome is associated with a much more severe recession. The peak-to-trough fall in aggregate consumption more than doubles, going from about 9 percent without containment measures to about 21 percent with containment measures. The mechanism underlying this

result is straightforward: higher containment rates make consumption more costly, so people cut back on the amount they consume and work.

Why not choose initial containment rates that are sufficiently high to induce an immediate, persistent decline in the number of infected? Absent vaccines, the only way to prevent a recurrence of the epidemic is for enough of the population to acquire immunity by becoming infected and recovering. The optimal way to reach this critical level of immunity is gradually increase containment measures as infections rise and slowly relaxing them as new infections wane.

5.1 Medical preparedness model

Comparing Figures 3 and 4 we see that the optimal containment policy is more aggressive in the medical preparedness model than in the SIR-macro model. The peak containment rate is a bit higher in the medical preparedness model (51.4 versus 47.3 percent) and occurs substantially earlier (at week 27 versus week 40). In addition, the containment rate comes down much more slowly in the medical preparedness model. These differences reflect that, other things equal, the social cost of the externality is much larger. Not only do agents not internalize the cost of consumption and work on infection rates, they also don't internalize the aggregate increase in mortality rates.

The optimal containment policy greatly reduces the peak level of infections from 3 to 1.2 percent reducing the death toll from 1.7 to 0.4 percent of the initial population. For a country like the U.S., this reduction represents roughly 4.3 million lives saved.

5.2 The treatment and vaccines models

Comparing Figures 3 and 5 we see that the optimal containment policy in the treatment and SIR-macro models are very similar. In the treatment model, along a path were no treatment is discovered, the optimal containment policy reduces the peak level of infections from 5.1 to 2.6 percent reducing the death toll from 0.53 to 0.36 percent of the initial population. This reduction represents roughly 0.6 million lives saved in the U.S.

The black-dashed lines in Figure 6 show that optimal policy is very different in the SIR-macro model and the vaccination model. With vaccines as a possibility, it is optimal to immediately introduce severe containment measures to minimizes deaths. Those containment measures cause a very large, persistent recession: consumption falls by about 14 percent for

roughly 50 weeks. But this recession is worth incurring in the hope that the vaccination arrives before many people get infected.

It is optimal to reduce and delay the peak of the infections in anticipation of a vaccine being discovered. Figure 6 displays the behavior of the vaccines model under optimal containment policy on a path where a vaccine does not arrive. Compared to the competitive equilibrium (solid blue lines), the peak of the infection rate drops from 5 percent to 2.5 percent of the initial population. Moreover, the infection peak occurs in period 51 rather than in period 33. Absent a vaccination arriving, the optimal policy reduces the death toll as a percent of the initial population from 0.53 percent to 0.45 percent. For the U.S. this reduction amounts to about a quarter of a million lives.

Above we discussed why it is not optimal to introduce immediate containment measures in the baseline SIR-macro and treatment models. But why is optimal policy so different in the vaccination model? The basic reason is that unlike treatment, a vaccine does not cure infected people. The expected arrival of a vaccine also reduces the importance of building up the fraction of the population that is immune to a level that prevents the recurrence of an epidemic.

6 Conclusion

We extend the canonical epidemiology model to study the interaction between economic decisions and epidemics. In our model, the epidemic generates both supply and demand effects on economic activity. These effects work in tandem to generate a large, persistent recession.

We abstract from many important real-world complications to highlight the basic economic forces at work during an epidemic. The central message of our analysis should be robust to allowing for those complications: there is an inevitable trade-off between the severity of the short-run recession caused by the epidemic and the health consequences of that epidemic. Dealing with this trade-off is a key challenge confronting policy makers.

Finally, we note that our model abstracts from various forces that might affect the long-run performance of the economy. These forces include bankruptcy costs, hysteresis effects from unemployment, and the destruction of supply-side chains. It is important to embody these forces in macroeconomic models of epidemics and study their positive and normative implications.

References

[1] Atkeson, Andrew "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios," manuscript, UCLA, March 2020.

[2] Farhi, Emmanuel and Ivan Werning "Dealing with the Trilemma: Optimal Capital Controls with Fixed Exchange Rates," NBER Working Paper No. 18199, June 2012.

[3] Julien Riou and Christian L. Althaus "Pattern of early human-to-human transmission of Wuhan," bioR$\not{x}$iv, January 24, 2020.

[4] Kermack, William Ogilvy, and Anderson G. McKendrick "A contribution to the mathematical theory of epidemics," Proceedings of the Royal Society of London, series A 115, no. 772 (1927): 700-721.

[5] Pueyo, Tomas "Coronavirus: Why You Must Act Now Politicians, Community Leaders and Business Leaders: What Should You Do and When?," *Medium*, March 10, 2020.

Appendix A Computing the Equilibrium

Guess sequences for $\{n_t^s, n_t^i, n_t^r\}_{t=0}^{H-1}$, assume a sequences for the exogenous variable $\{\mu_{ct}\}_{t=0}^{H-1}$ for some large horizon, H. In practice, we solve the model for $H = 150$ weeks. Compute the sequence of the remaining unknown variables in each of the following equilibrium equations:

$$\theta n_t^r = A\lambda_{bt}^r,$$

$$(c_t^r)^{-1} = (1+\mu_{ct})\lambda_{bt}^r,$$
$$u_t^r = \ln c_t^r - \frac{\theta}{2}\left(n_t^r\right)^2.$$

Iterate backwards to calculate:

$$U_t^r = u(c_t^r, n_t^r) + \beta U_{t+1}^r.$$

Proceed calculating the sequence for the one remaining unknown in the following equations:

$$(1+\mu_{ct})c_t^r = An_t^r + \Gamma_t \qquad (\lambda_{bt}^r),$$

$$\theta n_t^i = \phi A\lambda_{bt}^i,$$

$$\left(c_t^i\right)^{-1} = \lambda_{bt}^i,$$

$$u_t^i = \ln c_t^i - \frac{\theta}{2}\left(n_t^i\right)^2,$$

$$(1+\mu_{ct})c_t^s = An_t^s + \Gamma_t \qquad (\lambda_{bt}^s),$$

$$u_t^s = \ln c_t^s - \frac{\theta}{2}\left(n_t^s\right)^2.$$

Given initial values for I_1 and S_1, iterate the following six equations for $t = 0, .., H-1$:

$$T_t = \pi_{s1}(S_t c_t^s)\left(I_t c_t^i\right) + \pi_{s2}(S_t n_t^s)\left(I_t n_t^i\right) + \pi_{s3} S_t I_t,$$

$$Pop_t = S_t + I_t + R_t - D_t,$$

$$S_{t+1} = S_t - T_t,$$

$$I_{t+1} = I_t + T_t - (\pi_r + \pi_d) I_t,$$

$$R_{t+1} = R_t + \pi_r I_t,$$

$$D_{t+1} = D_t + \pi_d I_t.$$

Iterate backwards to calculate:

$$U_t^i = u(c_t^i, n_t^i) + \beta\left[(1-\pi_r-\pi_d) U_{t+1}^i + \pi_r U_{t+1}^r\right],$$

$$\tau_t = \frac{T_t}{S_t},$$

$$U_t^s = u(c_t^s, n_t^s) + \beta\left[(1-\tau_t) U_{t+1}^s + \tau_t U_{t+1}^i\right].$$

Proceed calculating the sequence for the one remaining unknown in the following equations:

$$\beta\left(U_{t+1}^i - U_{t+1}^s\right) - \lambda_{\tau t} = 0,$$

$$\left(c_t^s\right)^{-1} - \lambda_{bt}^s(1+\mu_{ct}) + \lambda_{\tau t}\pi_{s1}\left(I_t C_t^I\right) = 0.$$

Finally, use a gradient-based method to adjust the guesses $\{n_t^s, n_t^i, n_t^r\}_{t=0}^{H-1}$ such that the following three equations hold with arbitrary precision:

$$(1+\mu_{ct})c_t^i = \phi A n_t^i + \Gamma_t \qquad (\lambda_{bt}^i),$$

$$\mu_{ct}\left(S_t c_t^s + I_t c_t^i + R_t c_t^r\right) + \mu_{nt}\left(S_t n_t^s + I_t n_t^i + R_t n_t^r\right) = \Gamma_t,$$

$$-\theta n_t^s + A\lambda_{bt}^s + \lambda_{\tau t}\pi_{s2}\left(I_t n_t^I\right) = 0.$$

Figure 1: SIR-Macro Model vs. SIR Model

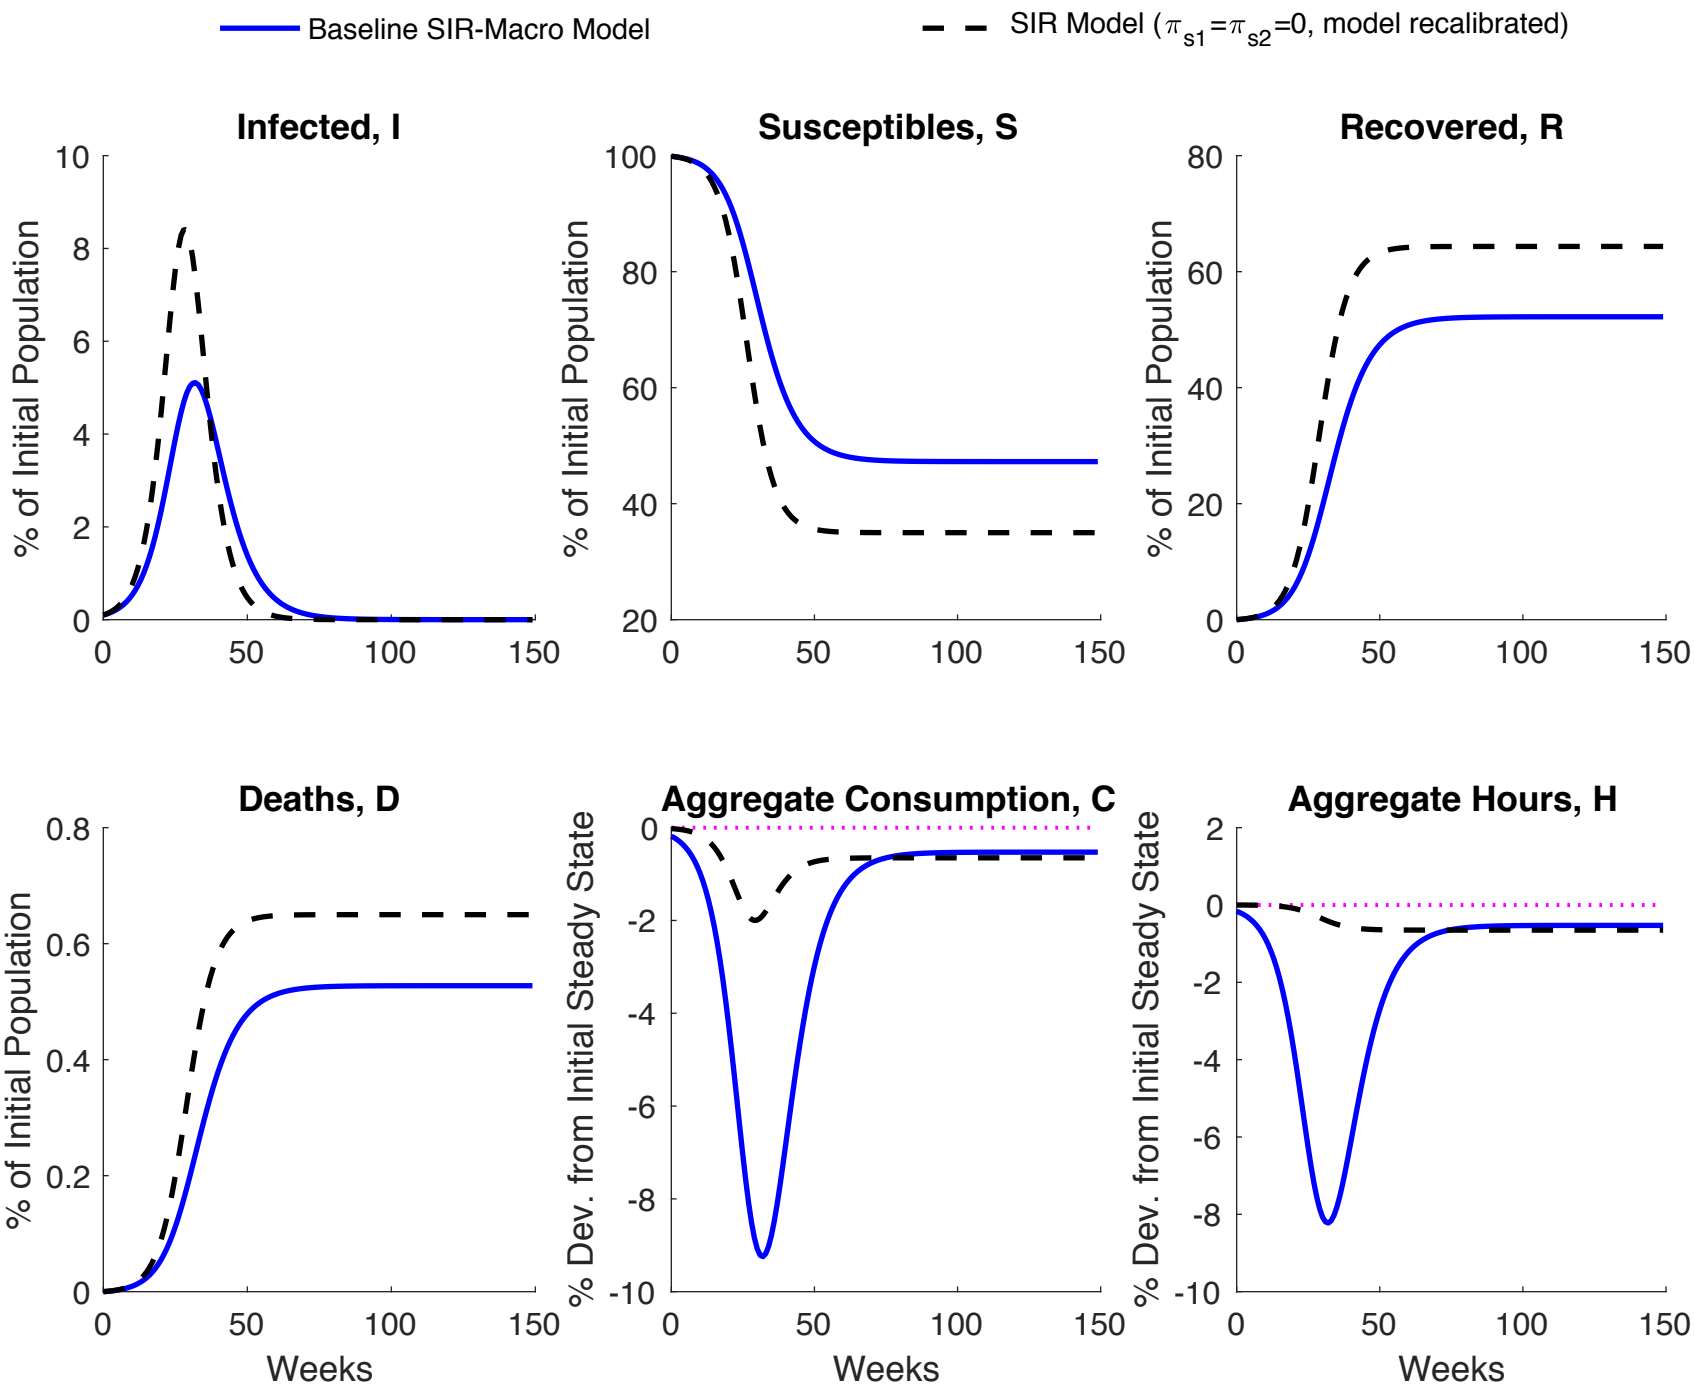

Figure 2: Consumption and Hours by Type in SIR-Macro Model

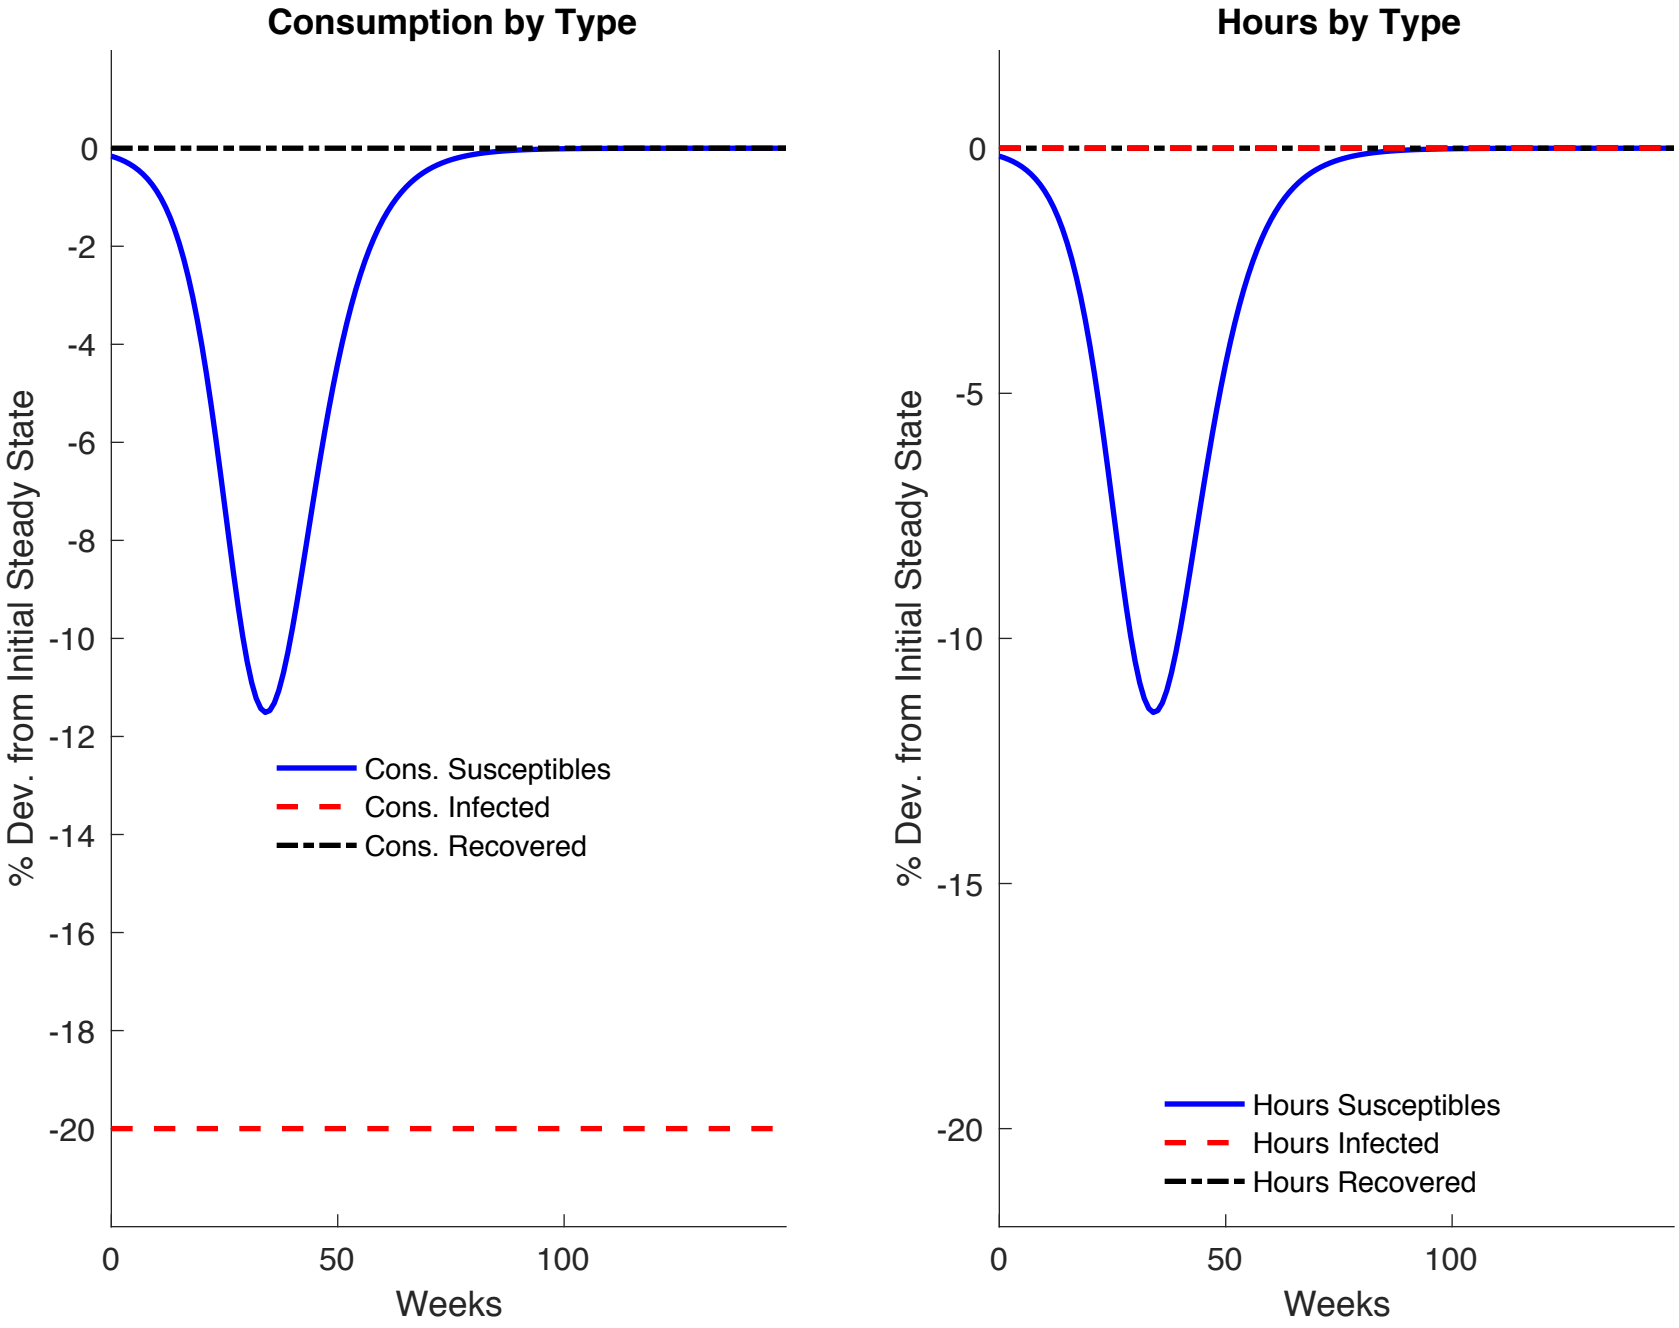

Figure 3: Baseline vs. Optimal Containment Policy

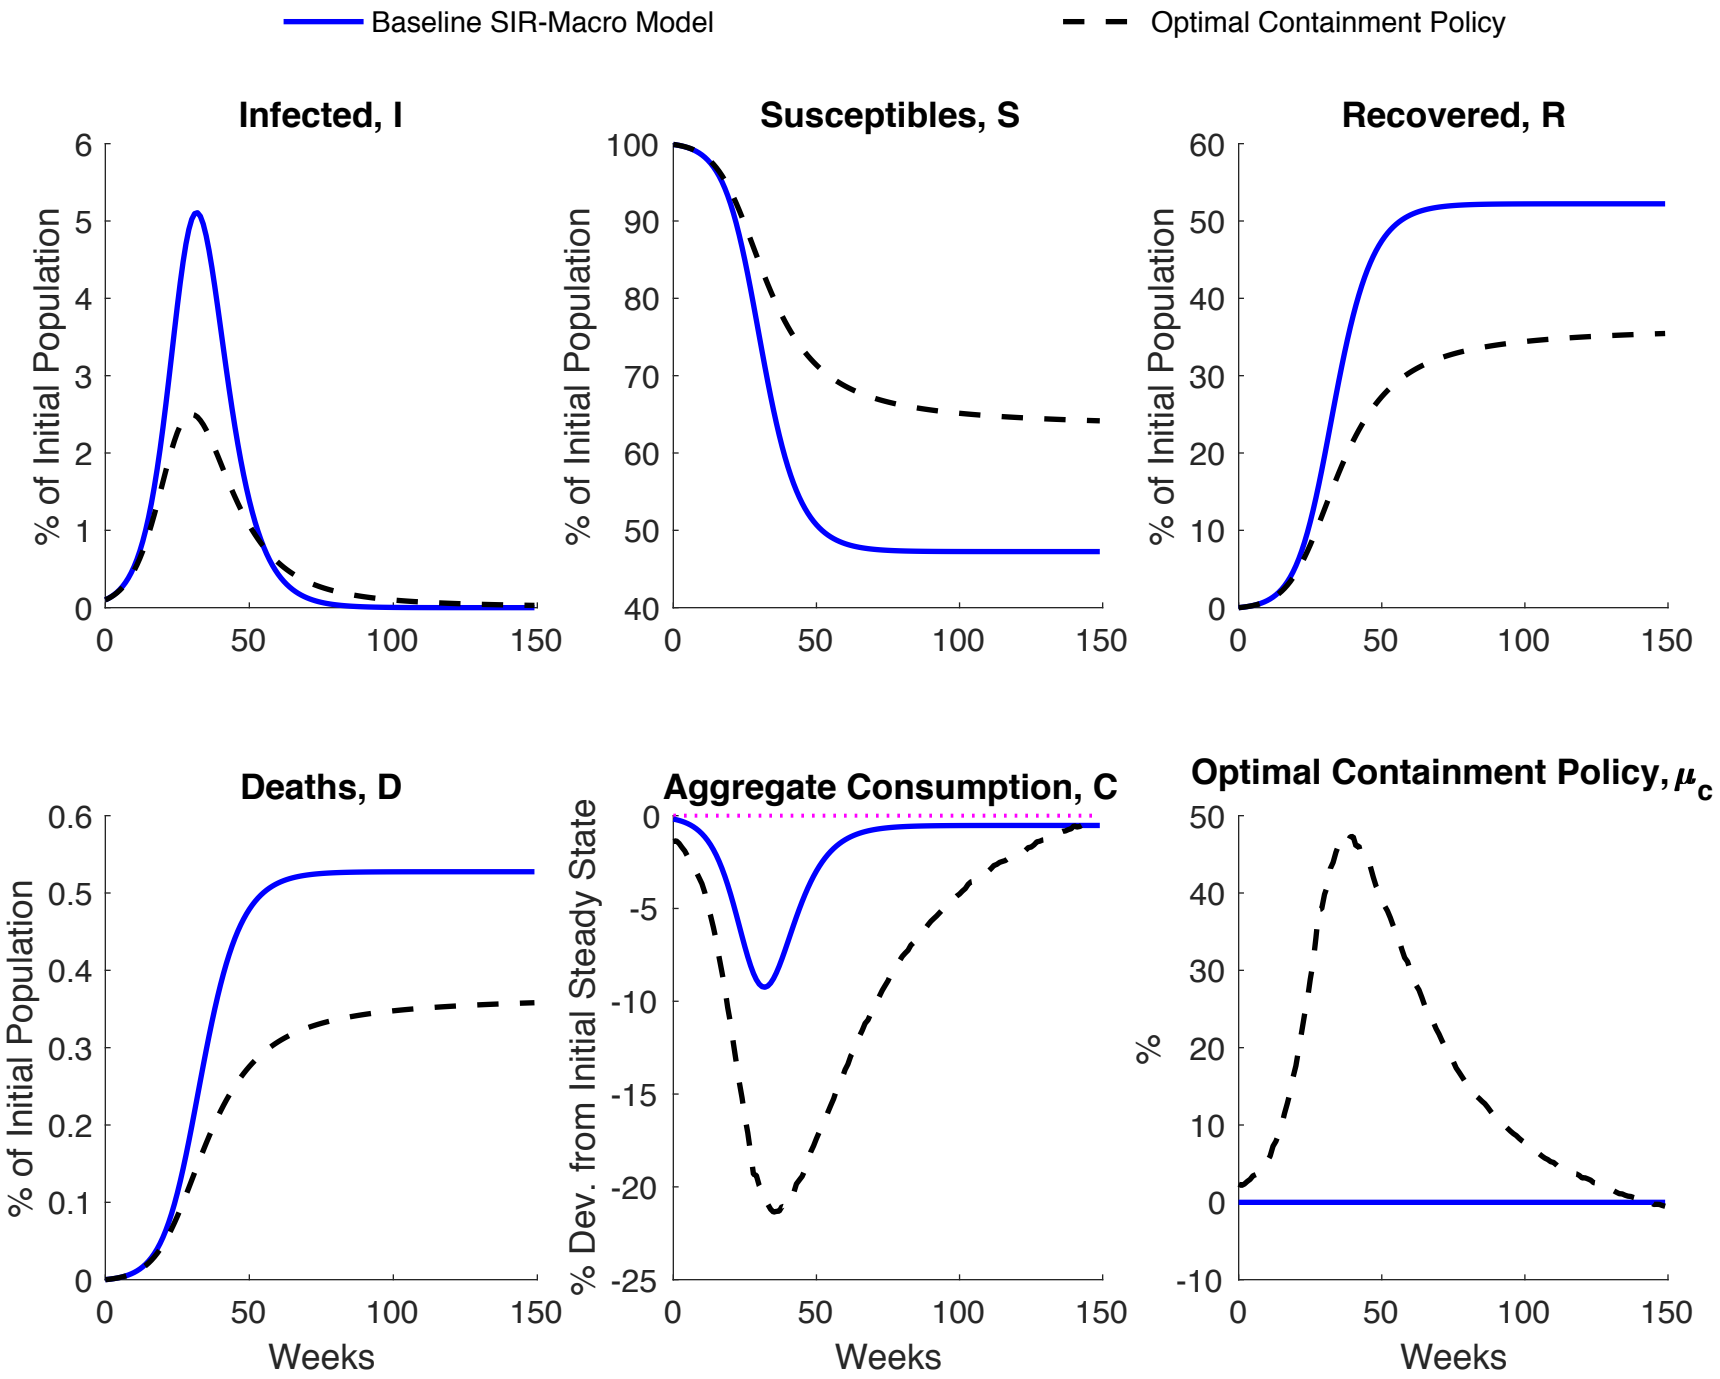

Figure 4: Medical Preparedness

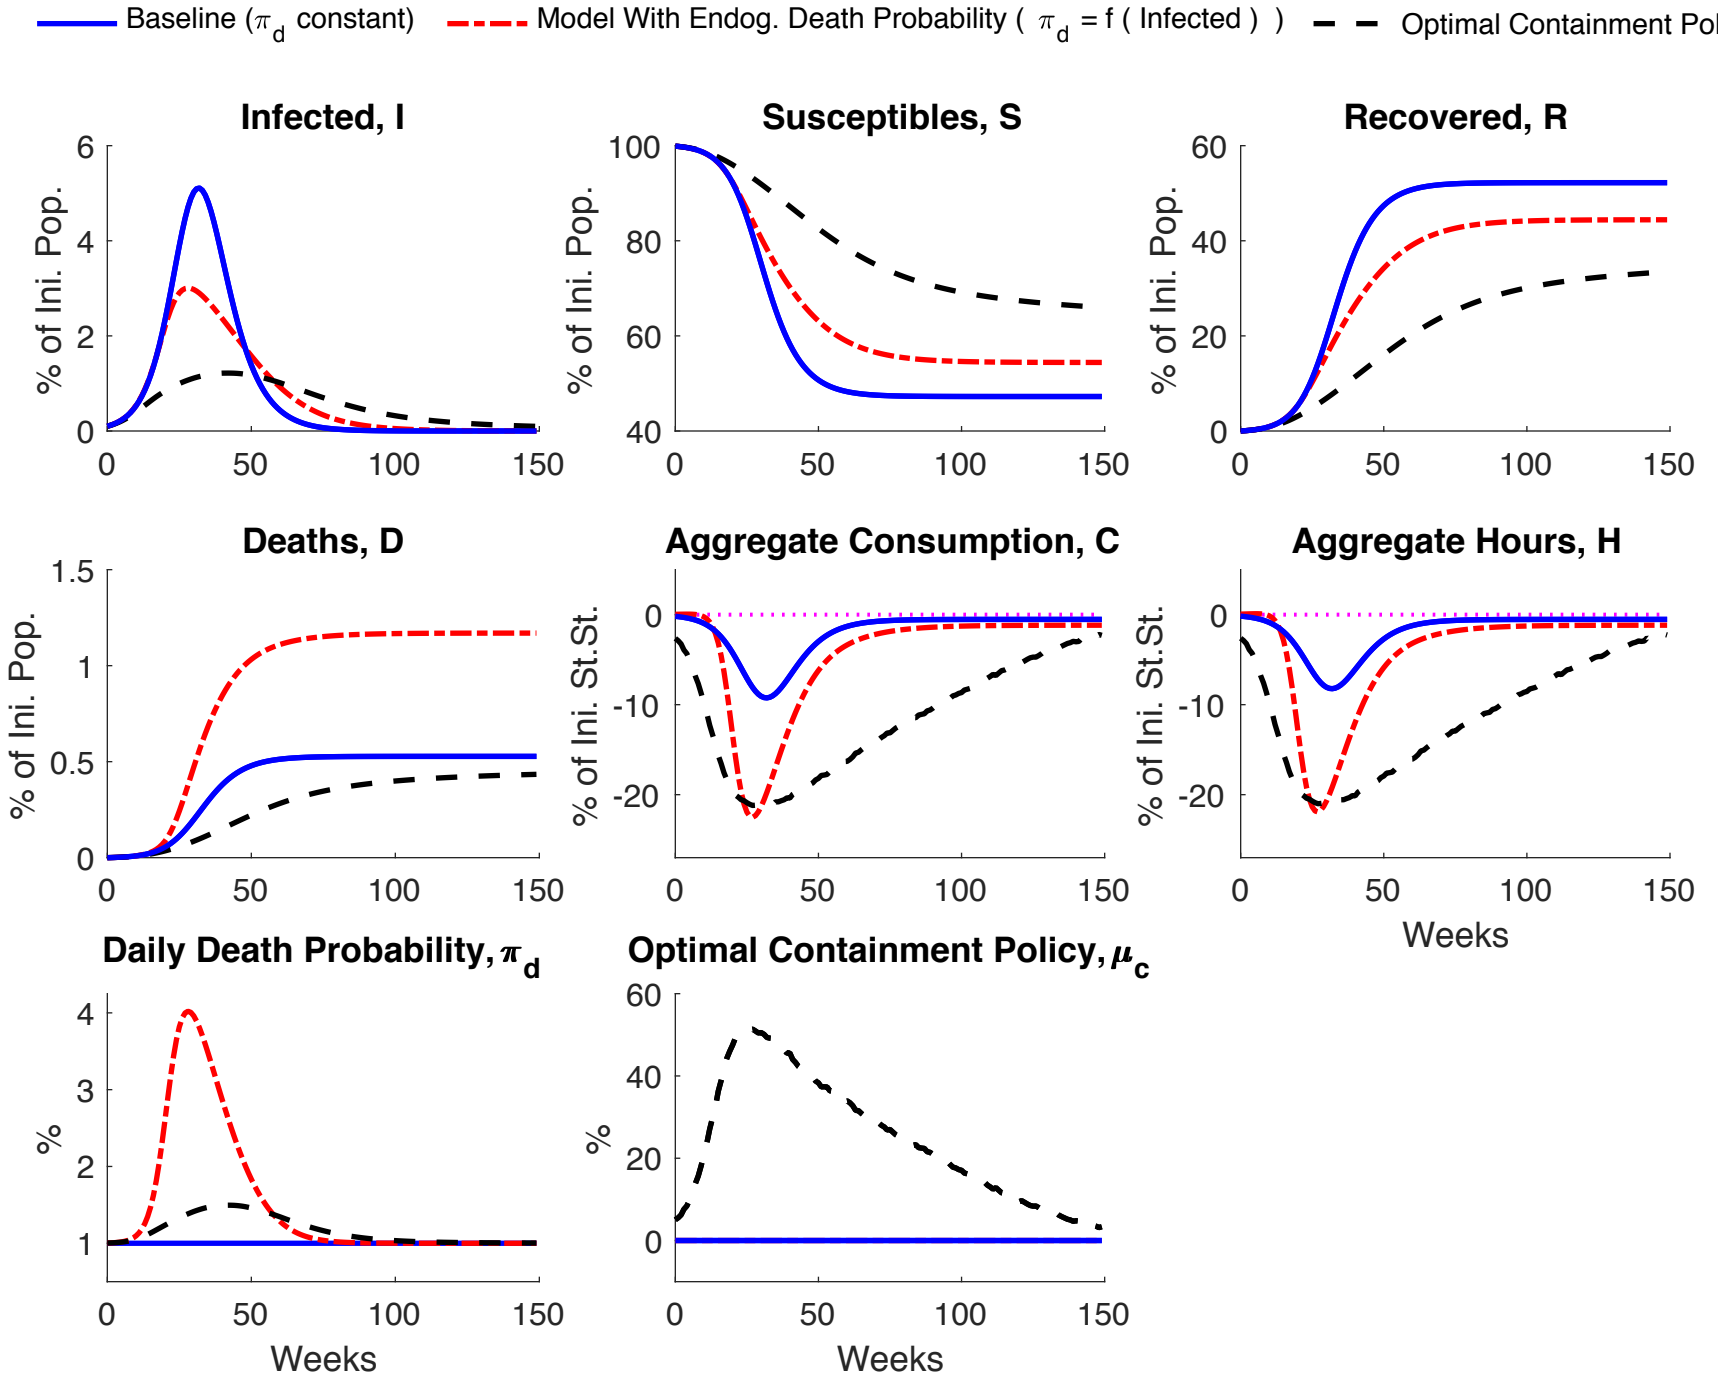

Figure 5: SIR-Macro Model With Treatments

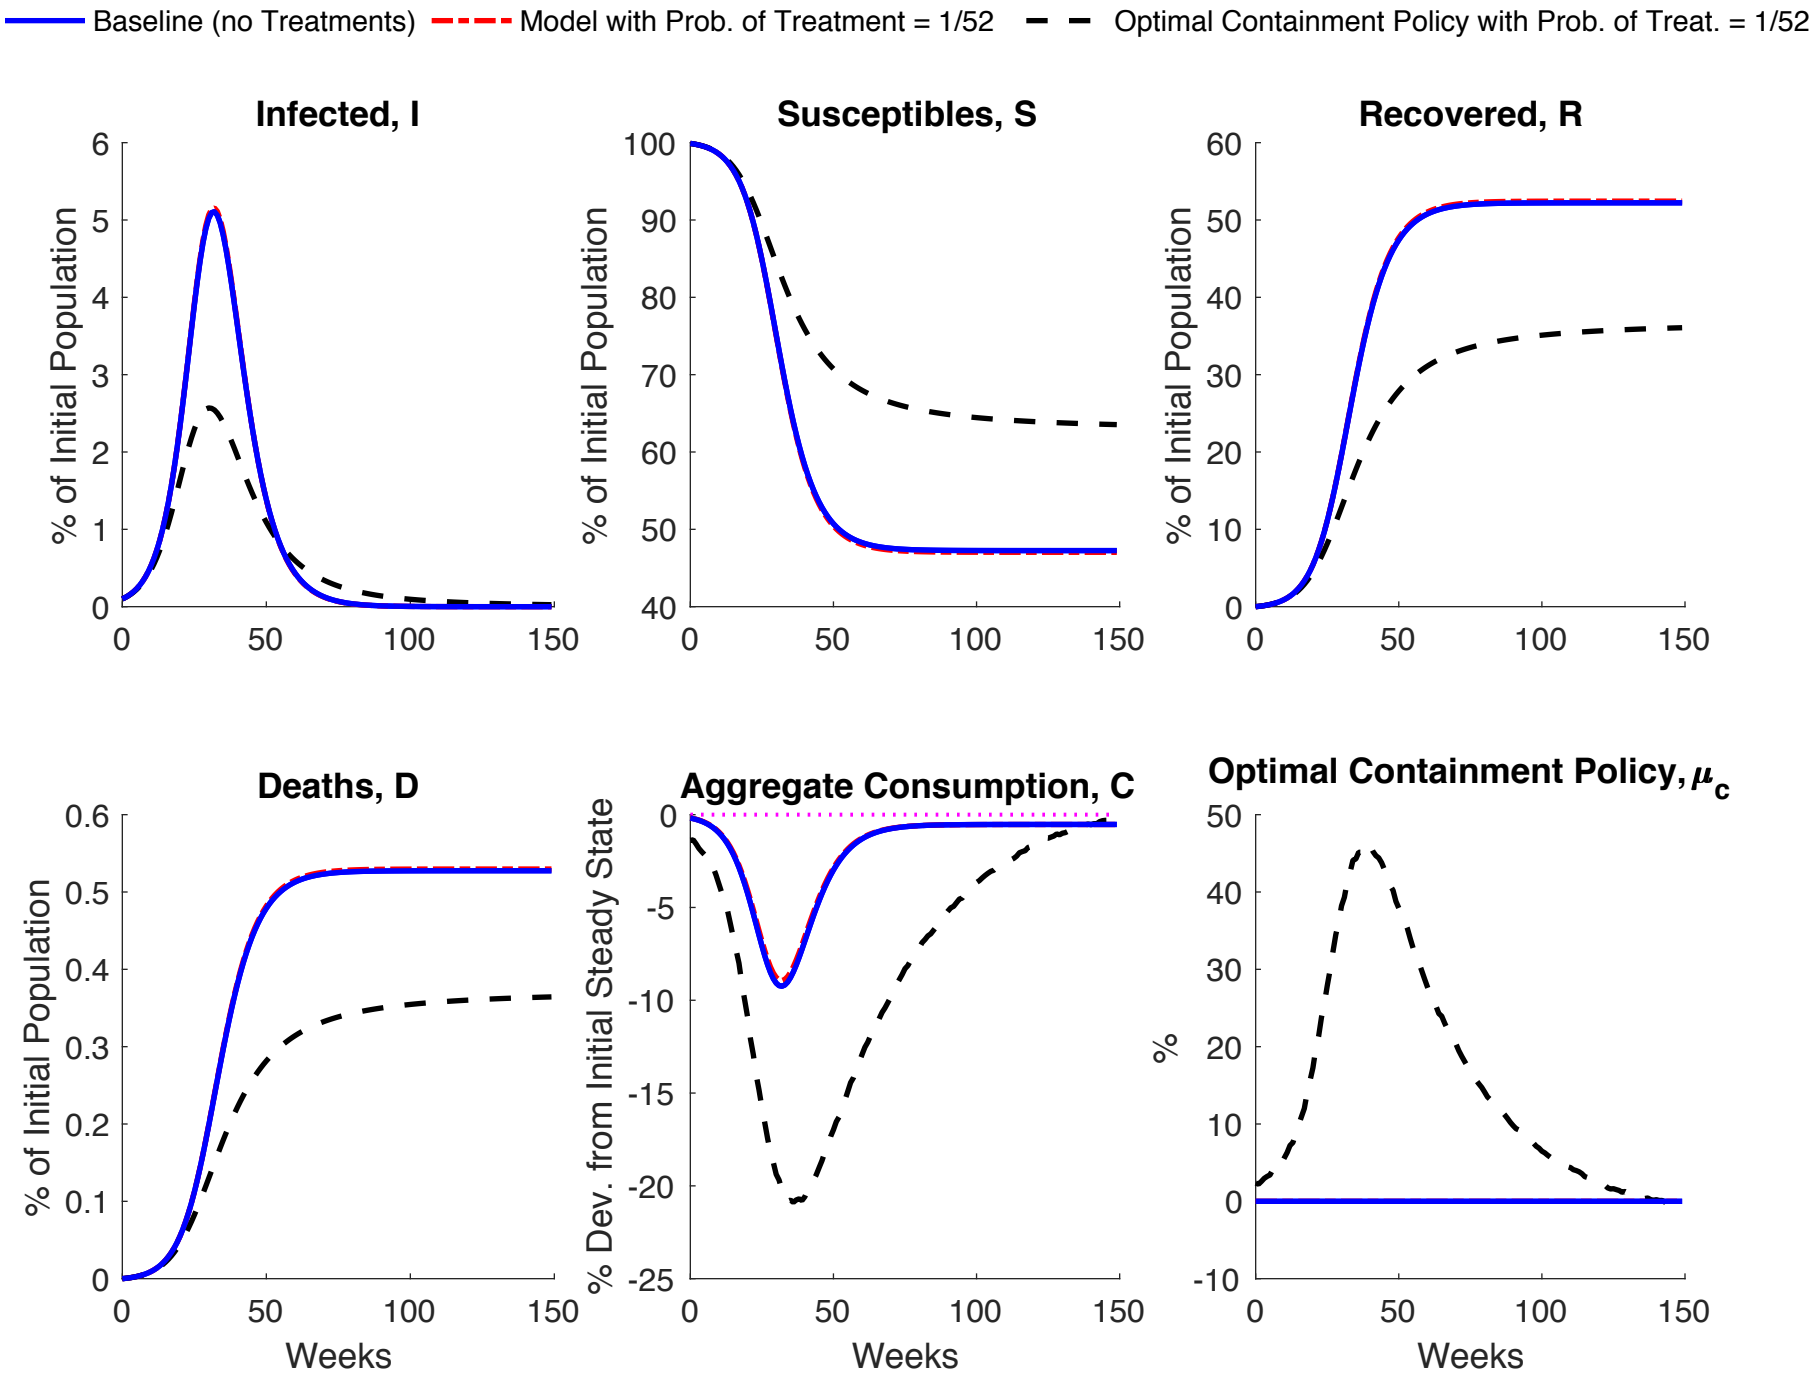

Figure 6: SIR-Macro Model With Vaccines