NBER WORKING PAPER SERIES

PANDEMICS, INCENTIVES, AND ECONOMIC POLICY: A DYNAMIC MODEL

Roberto Chang
Humberto Martínez
Andrés Velasco

Working Paper 28636
http://www.nber.org/papers/w28636

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2021

We thank Drew Harris and seminar participants at Rutgers University for useful comments. Any additional suggestions will be greatly appreciated. This project started while Chang was serving as BP Centennial Professor at LSE, whose hospitality he acknowledges with thanks. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Pandemics, Incentives, and Economic Policy: A Dynamic Model
Roberto Chang, Humberto Martínez, and Andrés Velasco
NBER Working Paper No. 28636
April 2021
JEL No. E6,F4,H3,I3

ABSTRACT

The advent of a pandemic is an exogenous shock, but the dynamics of contagion are very much endogenous --and depend on choices that individuals make in response to incentives. In such an episode, economic policy can make a difference not just by alleviating economic losses but also via incentives that affect the trajectory of the pandemic itself. We develop this idea in a dynamic equilibrium model of an economy subject to a pandemic. Just as in conventional SIR models, infection rates depend on how much time people spend at home versus working outside the home. But in our model, whether to go out to work is a decision made by individuals who trade off higher pay from working outside the home today versus a higher risk of infection and expected future economic and health-related losses. As a result, pandemic dynamics depend on factors that have no relevance in conventional models. In particular, expectations and forward-looking behavior are crucial and can result in multiplicity of equilibria with different levels of economic activity, infection, and deaths. The analysis yields novel policy lessons. For example, incentives embedded in a fiscal package resembling the U.S. CARES Act can result in two waves of infection.

Roberto Chang
Rutgers University
Department of Economics
75 Hamilton Street
New Brunswick, NJ 08901
and NBER
chang@econ.rutgers.edu

Humberto Martínez
Rutgers University
Department of Economics
75 Hamilton Street
New Brunswick, NJ 08901
hm409@economics.rutgers.edu

Andrés Velasco
School of Public Policy
London School of Economics
and Political Science
Houghton Street
London WC2A 2AE
United Kingdom
and CEPR
A.Velasco1@lse.ac.uk

Computer codes are available at https://github.com/totuma87/Chang_Martinez_Velasco_2021

1 Introduction

The Covid-19 pandemic has stimulated an avalanche of research from economists. Almost all this research treats the pandemic as an exogenous shock and then analyzes policy proposals to alleviate its consequences. But while the initial appearance of the virus may have been exogenous, its speed of propagation is not. The dynamics of the pandemic are the result of human decisions, and hence depend on economic policies.

Think of health policy directives such as lockdowns, stay-at-home guidelines, and social distancing rules. Their efficacy depends on the decisions of individuals who weigh the perceived costs against the benefits of compliance. A person who cannot work from home may be foregoing income by complying with a lockdown. But by staying home and reducing the chances of contagion that person also makes it more likely that she will be healthy and in a position to go to work when the economy opens up again. In turn these choices depend on a host of economic variables, including the size of foregone wages, the extent of government aid for those who stay at home, or expectations of the speed of the post-lockdown economic recovery.

In this paper we develop a dynamic model of the interaction between individual decisions and the speed of virus transmission, emphasizing the key role of economic incentives and the implications for policy. We extend the model of Chang and Velasco (2020), which assumed two periods, to an infinite time horizon. Aside from the obvious gain in realism, this change also makes the model easily comparable to the conventional Susceptible-Infected-Recovered (SIR) model of pandemics.

Just in SIR models, infection rates depend on human behavior. But here, unlike SIR models, behavior reflects the rational responses of individuals to incentives they face. Those incentives are given not only by the "technology" of infection but also by several financial costs associated with the pandemic. The interactions among virus transmission, incentives, and behavior have multiple and important implications. Incentives built into economic policy can affect the dynamics of virus infection. And while SIR models are purely backward-looking, expectations and forward-looking behavior are crucial in our model.

Our model features a population that normally works outside the home. Normality is interrupted by the appearance of a contagious virus. Infected individuals can become sick, in whose case they must stay in a "hospital" where lucky ones recover and unlucky ones die. People not showing symptoms of the virus choose whether to go to work or stay at home. The probability of infection is greater at work than at home, but staying at home is costly in terms of lost income. Hence agents decide between outside work and staying home on the basis of a "double relative": the current payoff from outside work relative to home income, compared against the expected future value of remaining asymptomatic relative to the cost of becoming ill with the virus.

In evaluating the double relative, individuals take into account the evolving probabilities of infection at home and at work. In turn, infection probabilities depend on the numbers of agents in each location. So there is a mutual interaction between location decisions and the evolution of the pandemic. In equilibrium, work-stay-home decisions and infection probabilities are simultaneously determined.

We are able to derive several analytical results. The model is highly nonlinear and time-dependent, however, so that resulting dynamics are generally complex and depend on parameters. To derive additional insights we simulate equilibrium paths numerically with parameters calibrated to match U.S. data.

For our benchmark calibrations, the model yields a pattern of infection with the hump shape characteristic of a pandemic. But because our model takes into account individual decisions and the corresponding incentives, equilibrium dynamics depend on fundamentals that would have no bearing in other models. For instance, increased risk aversion induces agents to spend more time at home during the pandemic, reducing infection rates and the prevalence illness and deaths. In the same vein, changes in the technology of infection affect equilibrium outcomes not only directly but also via the induced changes in individual behavior. For example, higher values for the parameter that determines contagion directly raise the speed of transmission. But, to avoid contagion individuals are more inclined to stay at home, which acts as an offsetting force.

More importantly for our discussion, changes in economic variables, which would not have an impact in SIR models, matter for the decisions of individuals and for the dynamics of virus transmission of our model. For instance, the relative rewards of staying at home versus outside work turn out to be key for dynamics. And those rewards, of course, can be affected by economic policies.

To emphasize this point, we examine the implications of a fiscal policy package similar to the U.S. 2020 CARES Act. We feed into the model an increase in unemployment benefits which virtually eliminates the difference between market income and home income, but only for a limited time period. In our model, such a policy causes the pandemic to come in two waves, consistent with the observed evolution of Covid-19 in the US.

The intuition is clear. While they last, unemployment benefits enacted by a policy such as the CARES Act reduce the relative payoff of working outside the home. These incentives prompt a change in behavior which helps limit contagion, contributing to a reduction in infections after an initial peak. However, the expiration of the CARES Act means that the payoff of outside work relative to staying home jumps back to its usual level. In response to the changed incentives individuals return to working outside the home, helping start a second wave of the pandemic.

Our model predicts that in the absence of the CARES Act the U.S. would have experienced a single peaked pandemic. That would also be the outcome if individual locations were exogenous, as in the SIR type of models. Furthermore, our analysis underscores that the parts of the CARES Act that matter for the evolution of the pandemic were those that affect individual incentives and decision-making. In contrast, provisions in the Act that do not change incentives have no impact on dynamics. This is the case, for example, of cash transfers to households regardless of whether they stay home or not, which in our model have no impact on relative payoffs and hence no effect on equilibria.

Our model is built on first principles, so that the analysis of several policies and their incentive effects is straightforward. As an illustration, "social distancing" can be thought of as a policy that reduces the number of contacts people have when they go to work. In our model, this policy has an ambiguous impact on behavior and contagion. Conditional on people's choices, fewer contacts outside the home reduce current and future infection probabilities. These changes have two opposite effects on individual incentives. A lower infection risk directly increases the relative payoff of working outside the home. But if future infection risks are lower as well, the opportunity cost of getting infected today and missing out on a brighter future can also go up, shifting incentives in favor of staying at home. So what is crucial for individual maximization is not the absolute value of infection rates but the "double relative" we explained earlier. In our calibrations, social distancing policies do have an ambiguous impact: depending on parameter values, the requirement that people stay socially distanced can increase the share of the population going to work for one set of parameter values, and reduce it for another.

The model incorporates an externality: individuals base their choices of location on their own probabilities of contagion, but fail to take into account the impact of those choices on the economy-wide dynamics of contagion. We contrast the decentralized outcome to the solution of the social planning problem in our model, one in which the social objective is the discounted expected welfare of the different groups, each weighted by their population size.

With the benchmark calibration, the planning solution at first restricts time outside the home drastically, and then allows agents to return to outside work gradually. The solution resembles a "lockdown" and stands in contrast to the decentralized equilibrium, in which people are fully working outside the home. The social-planning policy reduces infection but does not eliminate it completely.

Beyond comparing optimal versus decentralized outcomes, we ask: what is an individual's best response if the implements the optimal lockdown? For an individual who believes that others will comply with the lockdown, implying lower infection risk and weaker incentives to stay at home, the best response is to return to work full time. This suggests the optimal policy is politically very difficult to implement. Without compulsory enforcement or an economic policy that reduces the financial advantage of working outside the home, lockdown compliance will be, at best, short-lived.[1]

Finally, we analyze the possibility of multiple equilibria. *Animal spirits* can be crucial in a pandemic. For some parametrizations, a SIR-type equilibrium, with agents fully working outside the home and high rates of infection and death, coexists with another equilibrium in which precautionary behavior leads to less time outside the home, lower infection rates, and fewer deaths. In the precautionary equilibrium individual decision makers reduce their time working outside the home not because they face high infection rates if they go to work today, but because they expect the pandemic to subside enough over time so that the future value of not being infected is high. Multiple equilibria reflect the importance of forward-looking behavior, in sharp contrast to SIR models, which are only backwards-looking.

The existence of multiple equilibria implies that policy can be highly effective if government credibly acts as a coordinating device. If a SIR-type equilibrium coexists with a precautionary equilibrium, an appropriately tailored directive to limit non-essential time outside the home not only eliminates the bad equilibrium, but also is politically costless: in the surviving precautionary equilibrium, decision-makers find it individually optimal to follow the directive. In the same way, economic incentive policies can in principle act as coordinating mechanisms.

This paper is a contribution to the literature, motivated by the Covid-19 pandemic, that has attempted to blend models from economics and from epidemiology. The dominant epidemiological models of virus transmission are variants of the SIR model developed about a century ago (the earliest published version seems to be Kermack and McKendrick 1927). The paper by Weiss (2013) is an excellent presentation of the basic technical details of the model. Weiss also discusses how the SIR model provides the theoretical underpinning for stay-at-home directives, social distancing, and other public health policy responses to a pandemic. In the SIR model, such policies are the only game in town, because individuals act mechanically, with incentives playing no role. By contrast, our analysis highlights how policies, including economic policies,

[1] These results are consistent with Levy-Yeyati et al (2020) who empirically show that lockdown compliance declines with time, and is lower in countries with stricter quarantines, lower incomes and higher levels of labor precariousness

can shape incentives and individual behavior in a pandemic, which in turn affects the dynamics of virus transmission.

Much like the epidemiological literature, the recent economics literature related to Covid 19 has largely ignored the role of economic behavior and incentives in determining the trajectory of the pandemic.[2] There are exceptions, however, to which our paper is related. A number of studies, including Garibaldi, Moen, and Pissarides (2020), Rachel (2020), and Toxvaerd (2021), develop extensions of the SIR model in which each individual chooses a degree of social distancing, or something similar, which indirectly determines exposure to the virus. In turn, aggregate choices affect the SIR equations and the dynamics of infection. Because individual choices depend on perceived infection probabilities, these models also feature the kind of mutual interaction between individual decision-making and virus dynamics emphasized in our paper.

In these models and ours, externalities drive a wedge between the decentralized equilibrium and the social planning outcome. In existing contributions, the focus has been to characterize differences between these two potential outcomes, and to derive suggestions for lockdowns and other public health measures. By contrast, we emphasize the impact of economic policy incentives and of policy interventions such as the CARES Act.

In addition, we build the model from fundamental assumptions about technology, contacts, and infection technology, in a way that allows for a consistent analysis of policy questions and makes extensions easy to accommodate. Consider an example: several models in this group assume that infected individuals interact freely and knowingly with susceptible people. But if infection is observable, wouldn't susceptible agents simply avoid interacting with sick agents?. In our model, by contrast, we assume that people who become sick are sent to a hospital for treatment, so that contagion arises only from asymptomatic people who do not know they have been infected.

Eichenbaum, Rebelo, and Trabandt (2020) and Jones, Phillipon, and Venkateswaran (2020) develop dynamic models that postulate SIR-type equations, which depend on economic activit-

[2] Brodeur, Gray, Islam, and Jabeen Bhuiyan (2020) survey the economic literature as of June 2020. Since then, however, there have been many significant contributions.

ies such as consumption and hours worked. As in our paper, individual agents understand that their consumption and labor supply choices have implications for their exposure to contagion, so market incentives such as wages matter for the dynamics of infection. But our paper differs from Eichenbaum et al. and Jones et al. in several respects, some of which are significant for policy analysis. For example, both Eichenbaum et al. and Jones et al. assume that contagion increases with the levels of aggregate consumption. An implication is that raising consumption taxes during a pandemic would reduce infections, which would amount to an argument in favor of such a policy. In contrast, consumption taxes, because they wash out in the double relative, have no impact on individual choices in our model, and therefore no effects on contagion dynamics.

Whether these differences are important hinges on the specific objectives of the distinct papers. The main focus of Eichenbaum et al. and Jones et al. is on describing and quantifying dynamic implications. Given such a focus, there is no big loss in treating the link between consumption and infection rates as a reduced form. For policy analysis, on the other hand, the details of such a link are crucial, which is one of the reasons why we have placed special care in deriving SIR-type equations from first principles.

Our analysis of how fiscal policy affects the dynamics of a pandemic via incentives, and how fiscal policy can be a factor underlying multiple waves of infection, is new in the literature. Moreover, and to the best of our knowledge, we know of no model of a pandemic in which the interaction between economic incentives and forward-looking behavior can result in multiple equilibria. In turn, those different equilibria have very different implications for contagion dynamics and the associated policy implications.[3]

The rest of the paper proceeds as follows. Section 2 describes the economic environment in normal times. Section 3 describes the impact of a virus, emphasizing the interplay between individual decisions and infection rates, and defining equilibrium. Section 4 proposes a numerical calibration of the model and investigates model dynamics. We discuss the implications

[3] Multiple equilibria may emerge in SIR-type models of equilibrium social distancing. See Chen (2012).

of a fiscal package similar to the CARES Act in section 5. Social distancing is the focus of section 6 and the social planning problem is focus of the section 7, where we also compare this outcome with the decentralized equilibrium of the model. Multiple equilibria and the role of forward-looking behavior are discussed in section 8, while section 9 closes with conclusions and suggestions for further research.

2 A Basic Economy

Time is discrete and indexed by $t = 0, 1, 2....$The economy is populated by a continuum of agents. The size of the population is normalized to one.

People in this economy transit between three locations which we shall call *home*, *outside*, and *hospital*. As in Chang and Velasco (2020), transitions between locations depend partly on individual decisions, reflecting agents' perceptions of costs and benefits, including the possibility of infection.

There is only one final good. Each agent outside her home in period t receives an amount w_t. We can think of w_t as a wage or the amount of output that the agent produces working outside, and more generally a reward from market participation. Likewise, agents at home or in the hospital receive an amount of goods e_t. This can be thought of output from home production, or a subsidy from the government.[4] For now, we simply assume that w_t and e_t , $t = 0, 1, 2...$ are exogenously given, known sequences, with $w_t > e_t \geq 0$.

Agents consume their incomes in every period. In particular, we rule out borrowing or lending. This is in spirit of simplicity, but allowing for borrowing and lending may be a substantial extension.

In this economy, normal life is quite easy. Since $w_t > e_t$, agents spend every day working outside, and receive utility:

[4]We can allow the subsidy to differ between home and hospital, but here we assume there is no difference, for simplicity.

$$v_{zt} = \sum_{j=0}^{\infty} \beta^j u(w_{t+j}) \quad (1)$$

where $0 < \beta < 1$ is their subjective discount factor, and u displays constant relative risk aversion $\sigma > 0$:

$$\begin{aligned} u(c) &= \frac{c^{1-\sigma}}{1-\sigma} \text{ if } \sigma \neq 1, \\ &= \log(c) \text{ if } \sigma = 1 \end{aligned}$$

3 Pandemic

Things change, however, when, at the beginning of $t = 0$, a fraction $1 - h_0$ of the population gets infected with a virus. As we describe below, infected people show no symptoms until the end of the period. Hence, at the start of $t = 0$, people do not know if they are infected or not: they remain *vulnerable.* The number of vulnerable individuals at the start of any period t will be denoted by s_t. Hence $s_0 = 1$.

3.1 Vulnerables and Decision Makers

To describe the model dynamics, consider any period t, with s_t vulnerable individuals, of which a fraction $(1 - h_t)$ carry the virus. In other words, at the start of period t there are $(1 - h_t)s_t$ *asymptomatic* individuals, and $h_t s_t$ *healthy* ones.

For a benchmark case we select some simple assumptions, similar to those in Chang and Velasco (2020). An exogenous random fraction $q > 0$ of the vulnerables is selected and must go to work outside. One can think of this fraction as "essential" workers.

Each of the remaining $(1-q)s_t$ vulnerables decides what fraction of the period to stay home or outside. In the parlance of Chang and Velasco (2020), this is a group of *decision-makers.* We examine their decision problem shortly, but observe for now that a crucial consideration in that problem is that, in equilibrium, probabilities of infection are different in different locations.

We denote by ϕ_t^m (respectively ϕ_t^n) the probability that a *healthy* vulnerable gets infected if she spends a period working outside (resp. at home). Hence a healthy vulnerable that spends a fraction p_t of her day outside gets infected with probability

$$\bar{\phi}_t = \bar{\phi}_t(p_t) \equiv p_t\phi_t^m + (1 - p_t)\phi_t^n$$

In our model, the infection probabilities $\phi_t = (\phi_t^m, \phi_t^n)$ are endogenous, but are taken as given by individual agents.

Vulnerables do not know if they are healthy or infected at the beginning of the period (so each of them assumes that she is healthy with probability h_t). At the end of the period, however, some infected vulnerables become symptomatic. Let κ the fraction of infected that show symptoms. We assume that κ is exogenous and less than one.

Infected individuals that exhibit symptoms exit the vulnerable population to enter the hospital. We denote $x_{t+1}^{(1)}$ the number of vulnerables that are interned in the hospital at the end of period t:

$$x_{t+1}^{(1)} = \kappa\{(1 - h_t) + h_t[q\phi_t^m + (1 - q)\bar{\phi}_t(p_t)]\}s_t \tag{2}$$

In the preceding, the term in brackets is the fraction of vulnerables that are infected at the end of period t. That fraction is given by asymptomatic but infected vulnerables, $(1 - h_t)$, plus the number of healthy vulnerables that get infected during the period.

Correspondingly, the number of vulnerables next period is

$$s_{t+1} = s_t - x_{t+1}^{(1)} \tag{3}$$

and the number of healthy vulnerables in $t + 1$ is:

$$h_{t+1}s_{t+1} = h_t s_t \left\{1 - [q\phi_t^m + (1 - q)\bar{\phi}_t(p_t)]\right\} \tag{4}$$

3.2 Hospitalization and Recovery

Infected people that show symptoms recover only after spending some period of isolation or medical care in a hospital. Under our assumptions, people do not recover till they spend time in the hospital, hence they must show symptoms first. One interesting variation might be to allow asymptomatic people to recover without going to the hospital.

An individual in the hospital stays $H \geq 1$ periods there, after which she recovers with probability $(1-\mu)$ or dies with probability μ. A recovered person is virus free for ever, and able to earn the present value of outside wages, defined above as v_z. On the other hand, we assume that death involves a utility cost $D \geq 0$.

Hence the value of a hospitalized individual in her first day at the hospital is:

$$v_{ht} = \sum_{j=1}^{H} \beta^{j-1} u(e_{t+j-1}) + \beta^H \left[(1-\mu) v_{z,t+H} - \mu D\right] \tag{5}$$

Let $x_t^{(i)}$ denote the number of patients in their i^{th} day in the hospital, and $x_t = (x_t^{(1)}, ...x_t^{(H)})$. Also, let z_t denote the number of recovered people up to and including period t. Then $z_0 = 0$ and

$$z_{t+1} = z_t + (1-\mu) x_t^{(H)} \tag{6}$$

The law of motion of $x_t^{(1)}$ was given in (2). In turn, by definition,

$$x_{t+1}^{(i)} = x_t^{(i-1)}, \quad i = 2, 3, ...H \tag{7}$$

Finally, ω_t will denote the number of accumulated deaths. Then $\omega_0 = 0$ and

$$\omega_{t+1} = \omega_t + \mu x_t^{(H)}$$

Conditional on $\{p_t\}$ and $\{\phi_t\}$, the equations defined so far determine the evolution of

s_t, h_t, x_t, ω_t and z_t. In fact, as the reader may recognize, the equations are similar to those of the SIR model of virus transmission. But here $\{p_t\}$ and $\{\phi_t\}$ are not fixed parameters but equilibrium objects, determined by the decisions of agents in the model. We turn to this aspect of the model.

3.3 Individual Decisions

As mentioned, at the beginning of each period t, a fraction q of the vulnerables are exogenously sent outside. These agents do not have any decision to make. The value of their lifetime utility from then on, that is, their value function v_{qt}, is easily seen to be:

$$\begin{aligned} v_{qt} &= u(w_t) + \beta[\kappa\{(1-h_t) + h_t\phi_t^m\}v_{ht+1} \\ &\quad +(1-\kappa\{(1-h_t) + h_t\phi_t^m\})v_{st+1}] \end{aligned} \tag{8}$$

where v_{ht} is the value function at the hospital, and v_{st} is the value function for a vulnerable at time t, to be defined below.

The right hand side is the utility of the outside wage plus the discounted expected value of their utility from next period on. For the latter, observe that the probability of being sick at the end of the period equals the probability of being sick at the beginning of the period, $(1-h_t)$, plus the probability of starting healthy but infected outside during the period, $h_t\phi_t^m$. Also, a fraction κ of the sick population at the end of period becomes symptomatic and must exit to the hospital. Otherwise, the agent remains in the vulnerable population.

The remaining $(1-q)s_t$ vulnerables face the more delicate choice of how to distribute her time in or out of their homes. Crucially, this choice determines not only their current income but also their infection probabilities. Each agent in this group knows that, if she is healthy and spends a fraction p_t of the period outside, she will get infected with probability $\bar{\phi}_t = p_t\phi_t^m + (1-p_t)\phi_t^n$.

Hence the problem of such decision makers can be written as:

$$
\begin{aligned}
v_{dt} &= Max_{0\leq p_t \leq 1}\ u(c_t) + \beta\{\kappa[(1-h_t) + h_t\bar{\phi}_t]v_{h,t+1} \\
&\quad +[1-\kappa((1-h_t)+h_t\bar{\phi}_t)]v_{s,t+1}\} \\
\text{where } c_t &= p_t w_t + (1-p_t)e_t \\
\text{and } \bar{\phi}_t &= p_t\phi_t^m + (1-p_t)\phi_t^n
\end{aligned}
\tag{9}
$$

The maximand in the RHS reflects that, if the individual spends a fraction p_t of the period outside her home, her current income and consumption is $c_t = p_t w_t + (1-p_t)e_t$. In addition, she will be infected at the end of the period with probability $(1-h_t) + h_t\bar{\phi}_t$, in whose case she will become symptomatic with probability κ and enter the hospital next period, receiving $v_{h,t+1}$. Otherwise, she will remain in the vulnerable group, receiving $v_{s,t+1}$.

It is instructive to examine the derivative of the maximand with respect to the choice p_t, which is

$$
u'(c_t)(w_t - e_t) - \beta\kappa h_t(\phi_t^m - \phi_t^n)(v_{st+1} - v_{ht+1}) \tag{10}
$$

This is an illuminating expression. The first term is the current gain of a marginal increase in p_t. Such an increase raises current income by the utility value of outside income *relative to* home income, $w_t - e_t$. That gain is compared against the marginal cost associated with infection risk. What is that risk? By working outside rather than staying home, a vulnerable individual raises the chance that she is healthy but gets an infection. That increase is captured by $h_t(\phi_t^m - \phi_t^n)$. With probability κ the individual will then show symptoms and have to enter a hospital next period, with cost $\beta(v_{st+1} - v_{ht+1})$.

Hence decisions to work outside or stay at home depend on a "double relative": the current payoff to outside work relative to staying home is compared with the expected discounted value of future payoff of remaining vulnerable versus going to hospital. The choice problem has an intratemporal and an intertemporal dimension. Crucially, expectations play a crucial role.

Finally, our assumptions imply that

$$v_{st} = qv_{qt} + (1 - q)v_{dt} \tag{11}$$

And, of course, the decision problem depends on the probabilities of contagion, ϕ_t^m and ϕ_t^n. This will depend on the "technology" of virus transmission.

3.4 Contagion

We impose SIR-type assumptions, deriving contagion probabilities from basic assumptions about frequency of meetings and rates of transmission in different locations. In this way, one can think about a variety of policies, such as "social distancing", in a useful way.

Each agent outside her home has ρ^m close meetings with other agents during a period. A healthy agent contracts the virus with probability γ if she meets an infected person. In turn, the probability of meeting an infected individual in a given match is equal to the proportion of healthy agents outside, given by:

$$h_t^w = \frac{[q + (1 - q)p_t]h_t s_t + z_t}{[q + (1 - q)p_t]s_t + z_t} \tag{12}$$

taking into account that z_t recovered agents have returned to outside work and are healthy.

It follows that the probability that a healthy vulnerable agent working outside is *not* infected in a given meeting is $h_t^w + (1 - \gamma)(1 - h_t^w)$ and hence[5]

$$\phi_t^m = 1 - [h_t^w + (1 - \gamma)(1 - h_t^w)]^{\rho^m} \tag{13}$$

The expression is intuitive. An increase in the proportion of infected people in the market raises ϕ_t^m. Given h_t^w, an increase in the number of meetings, ρ^m, leads to an increase in ϕ_t^m. "Social distancing" policies are, presumably, those that attempt to reduce ρ^m. Finally, policies

[5] Note that CV's assumption is $\rho = \gamma = 1$.

such as mandating the use of face masks may affect the probability of transmission γ.

Analogous reasoning implies that

$$\phi_t^n = 1 - [h_t + (1-\gamma)(1-h_t)]^{\rho^n} \tag{14}$$

where ρ^n indicates the number of close meetings at home. It is natural to assume that $\rho^n < \rho^m$.

This completes the description of the model. Importantly for our purposes, contagion probabilities depend not only on the extent of infection but also on individual decisions, here given by how vulnerable agents allocate their time between outside work and home. But those decisions, as we have seen, depend on those same agents' perceptions of contagion probabilities.

3.5 Equilibrium

Equilibrium can be defined in a natural way. The definition includes the dynamics of infection, in addition to individual decisions.

A (**perfect foresight**) **equilibrium** involves sequences of population fractions $\{h_t, s_t, x_t, z_t\}$, value functions $v_{st}, v_{qt}, v_{dt}, v_{ht}$, and v_{zt}, time allocation decisions p_t, and contagion probabilities $\phi_t = (\phi_t^m, \phi_t^n)$ such that:

- Given $\{\phi_t, p_t\}$, $s_0 = 1, z_0 = 0$, and $x_0^{(1)} = ... = x_0^{(H)} = 0$, and a given $h_0 \in (0,1)$, $\{h_t, s_t, x_t, z_t\}_{t=1}^{\infty}$ satisfy 4, 3, 7, 2, and 6
- The value functions satisfy 8, 9,11, and 1, 5, given $\{\phi_t, h_t\}$
- p_t attains the max in the RHS of the Bellman equation 9
- ϕ_t^m and ϕ_t^n are given by 13 and 14

As we have emphasized, our model is similar to existing SIR models, except (crucially) that p_t is endogenous. This similarity allows us to immediately derive some qualitative features of equilibria, adapting arguments of e.g. Weiss (2013).

Assume that w_t and e_t are eventually constant. Also, consider a sequence $\{p_t\}$ such that $p_t = 1$ for all t sufficiently large (this will be an equilibrium feature). Then 3 and 2 imply that $\{s_t\}$ is a decreasing sequence bounded below by zero, so it must converge to some limit that we denote by $s^\infty \in [0,1]$. Likewise, $\{z_t\}$ and $\{\omega_t\}$ are increasing bounded sequences, so it must converge to some $z^\infty, \omega^\infty \in [0,1]$. It also follows that $\{x_t\}$ converges to the zero vector, and that $z^\infty + \omega^\infty = 1 - s^\infty$.

Hence, in the long run, the pandemic subsides. But does everybody get infected? Indeed, this can be the case under some parameter values that imply that $z^\infty + \omega^\infty = 1$. In such a case, everyone in the population gets infected, eventually goes to the hospital, and recovers or dies.

But it is also possible that $z^\infty + \omega^\infty < 1$ and $s^\infty > 0$. To see how, note that if h_t and h_t^w converge to one, the probabilities of infection ϕ_t^m and ϕ_t^n fall to zero. If this convergence is sufficiently fast, infections fizzle out while there is still a positive mass of vulnerables. This is a case of "herd immunity".

More detailed dynamics can be inferred by focusing on the number of new infections in each period, given by:

$$N_t = s_t h_t \left\{ [q + (1-q)p_t]\phi_t^m + (1-q)(1-p_t)\phi_t^n \right\}$$

that is, the number of initially healthy vulnerables at the start of the period, $s_t h_t$, times the probability that each of them is infected during the period. The number of healthy vulnerables is decreasing, but the probability of infection can increase or decrease. Consequently, N_t can increase or decrease, although eventually it must converge to zero. In typical SIR models, if $N_1 < N_0$, the convergence is monotonic, while if $N_1 > N_0$ there must be at least one "peak" in infections. This depends on the particular parameters of the model.

At this point, we cannot extract more implications of the model analytically. This is due to the high nonlinearity and nonstationarity of the model. But we can learn much more from the model by numerical study of specific parametrizations. We turn to these.

4 A Benchmark Case

4.1 Calibration

We assume that each time period of the model represents a day. Consistent with an annual interest rate of one percent, we set β equal to $(1/1.01)^{1/365}$.

We calibrate the model to the U.S. According to official numbers, by the end of the first week of March, there were about 338 active cases identified in the country.[6] However, it is well known that due to limited testing, the true number of active cases by that time could have been 10 to 25 times larger. Taking this into account, we set the initial fraction of healthy vulnerable population ($s_0 h_0$) equal to $1 - 10^{-5}$.

Parameters (γ, ρ^m, ρ^n) determine the rate of transmission of the virus outside and at home. Mossong et al. (2008) conducted a population-based prospective survey of mixing patterns in eight European countries using a common paper-diary methodology. They find that, on average, a person in a household of two to three people has between 10.65 and 12.87 daily contacts, of which 23% occur inside the household. Thus, we set the total average number of daily contacts to 11.7, which is consistent with the average household size in the US,[7] that implies a distribution of contacts between outside and home equal to $\rho^m = 9$ and $\rho^n = 2.7$. We set the probability of transmission per contact (γ) at 5% that lies in between 4.1% and 6.2%, which are the probability of transmission per contact with a asymptomatic and a symptomatic, respectively, estimated by He et al. (2020).

Recall that κ represents the fraction of infected that show symptoms at the end of each period. This parameter could also be interpret as the probability of showing symptoms. A joint mission by the World Health Organization and the Chinese government established that, on average, infected people developed signs and symptoms between 5 and 6 days after infection.[8] Consistently with this finding, we set $\kappa = \frac{1}{5.5}$.

[6] The CDC - Centers for Disease Control and Protection - see here
[7] According to the US Census, the average household size is 2.53 people
[8] See Page 11, Final Report of the mission here

Related to the virus, there are two additional parameters that need to chosen. The first is the number of days that a person spends at the hospital (H). Similar to existing literature on this matter,[9] we assume H is equal to 18, implying that after 18 days at the hospital, a person recovers or dies. We set the probability of dying (μ) to be 1% which is consistent with Verity et al. (2020) who estimated that the infection fatality rate in China, with a 95% of confidence, is between 0.39% -1.33%.

For the benchmark scenario, we set the economic parameters of the model as follows. We normalize w equal to one and choose $e_t = e = 0.38$. Thus, in the baseline scenario, the benefit of staying at home is 38% of the reward from working outside This number was chosen to reflect that in the pre-pandemic U.S., average national unemployment weekly payment was \$370 compared to the \$970 average national weekly salary of potential unemployment benefits recipients.[10] Parameter q is the probability that in each period, a vulnerable is selected to work outside. Following Alvarez et al. (2020), 30% of US GDP is generated by essential sectors, hence, we set $q = 0.3$ to capture that this fraction of the economy needs to operate in every period.

Finally, a key parameter for the calibration is D that corresponds to the utility loss upon death. How to calibrate this parameter is controversial. In the model, an obvious cost of death is the loss of wages. But it is easy to argue that it should include not only foregone earnings but also physical pain and suffering, and perhaps other considerations. Hence, for a benchmark, we take a pragmatic approach as follows. Kniesner and Viscusi (2019) indicate estimates of the value of a statistical life (VSL) for the U.S. are close to \$10 million (\$2017). We take this number to represent the expected present value of all costs associated with death, including not only wages, but also the additional costs just mentioned. We express those costs as a daily quantity, express that quantity as a constant times the daily average wage, and then compute D as the discounted value of the utility of the resulting constant. In the benchmark calibration,

[9] See Acemoglu et al. (2020), Eichembaum, Rebello, et al. (2020), Alvarez et al. (2020), and Verity et al. (2020) for example

[10] See Five Thirty Eight

we assume that agents have log utility preferences.

4.2 Equilibrium, Incentives, and the SIR Model

Figure 1 displays the predictions of the model during the first 150 days of the pandemic, assuming logarithmic utility.[11] The upper left hand panel shows that all agents spend all of their time outside their homes, in spite of the fact that the probability of contagion is lower at home. As a consequence, in this case the model behaves just as if we had assumed that agents had no choice between working outside and staying home. In other words, the model effectively becomes a standard SIR model.

Like in the SIR model, the pandemic results in a peak in infections at about eighty days from the initial seed. As vulnerables get infected, they transit to the hospital, where they either recover or die. Eventually the pandemic subsides. This reflects that the number of healthy vulnerables, susceptible to contagion, falls as more people acquire the virus. Also, hospitalized agents that recover return to work outside, increasing the relative number of healthy people there. In the long run, about one percent of the population dies.

Figure 1 confirms that our model delivers dynamics not unlike the SIR model. At the same time, however, it may give the misleading impression that, as in the SIR model, incentives are irrelevant. That this is not the case is illustrated by Figure 2, which compares the SIR model against ours in the case of a CRRA σ equal to ten.

The case $\sigma = 10$ is displayed by the red dashed line. As in the SIR model, the rate of infection starts accelerating about thirty days after the initial seed. Unlike in the SIR model, and in the case with log utility, the upper left panel of Figure 2 shows that decision making vulnerables start reducing their time outside, and more so as the infection rate goes up. About two months into the pandemic, the infection rate peaks, and the fraction of time outside bottoms at around one half. These two variables interact: since people stay home, the infection rate peaks at a lower level than in the SIR model. As a consequence, the number of active cases

[11] Figures are collected at the end of the paper. Computer code is available at ht-tps://github.com/totuma87/Chang_Martinez_Velasco_2021

falls, and the total number of deaths is lower, at eight tenths of one percent of the population, than in the SIR case (and also the log utility case).

Of course, the difference between Figures 1 and 2 amounts to an assumption about individual behavior. The SIR model features no decision making, while our model places it at center stage. The log utility case shows, however, that allowing for individual decision making is not sufficient by itself to depart from the SIR paradigm. People's responses to incentives must be strong enough, as in the case $\sigma = 10$.

Figure 3 illustrates the role of incentives. The upper left hand panel displays the marginal value of fully working outside given by the derivative of the objective function in the Bellman equation, (10), evaluated at $p_t = 1$. As shown, the derivative falls below zero after about thirty five days, expressing that fully working outside is suboptimal, so that decision makers increase time spent at home.

The upper right hand panel of Figure 3 shows the evolution of the term $\beta\kappa h_t(\phi_t^m - \phi_t^n)(v_{st+1} - v_{ht+1})$ in 10. As discussed before, this term captures the cost to decision makers of increasing outside working time p_t, due to the impact on infection risk. As the panel shows, it is the evolution of this term which explains the changes in incentives during the course of the pandemic. The term, in turn, reflects the differential infection risk outside vis a vis home, $h_t(\phi_t^m - \phi_t^n)$, which is displayed in the lower middle panel. But it also reflects the changing relative value of not being hospitalized, $\beta\kappa(v_{st+1} - v_{ht+1})$, displayed in the lower right panel.

To underscore the crucial role of individual responses to incentives, Figure 4 compares the original calibration, including log utility, against a case in which the utility cost of death is ten times larger. In our model, this means that vulnerable decision makers are more prone to stay home, rather than work outside, in order to reduce their probability of acquiring the virus and, ultimately, of death. Of course, the utility cost of death is irrelevant in the SIR model.

The results are intuitive. When people have a bigger fear of death, they choose to stay at home nearly one hundred percent of their available time as soon as infection rates start going up. As a consequence, infection rates and the number of active cases are much smaller than in the SIR case. The number of deaths falls to six tenths of one percent of the population.

5 Application: The CARES Act

5.1 The CARES Act

In the last days of March of 2020, the U.S. Congress reached an agreement to provide a economic relief package worth approximately \$2 trillion dollars. This package, denominated The Coronavirus Aid, Relief, and Economic Security Act (or CARES Act) came as a response to the economic fallout of the COVID-19 pandemic in the United States.

The CARES Act included the extension of unemployment benefits. The Federal Pandemic Unemployment Compensation (FPUC) provided an additional \$600 dollars per week to people for those receiving unemployment benefits. Additionally, the Pandemic Emergency Unemployment Compensation (PEUC) increased by 13 weeks the time window to receive unemployment benefits while the Pandemic Unemployment Assistance (PUA) expanded the eligibility criteria to self-employed and gig-workers.

To infer the effects of the CARES Act in our model, we focus on the impact of the Act on the relative compensation of working in the market versus staying at home. Relative to the SIR/log utility scenario, all parameters remain the same except for the schedule of the benefits of staying home, which we modify to capture the impact of the CARES Act on unemployment benefits.

Given that this policy was active between April and July, we assume that the implementation period starts 30 days after the onset of the pandemic and it lasts for an additional 120 days. Moreover, during this time frame, e_t increases from 0.38 to 0.99 capturing, as argued by Ganong et al. (2020), that the CARES Act increased the replacement rate for 76% of eligible workers above 100%. That is, most people would earn more from being unemployed than by working during this period of time. After day 150, e_t returns to the pre-pandemic value of 0.38.

Because of the short run nature of the problem at hand, we do not ask how the policy under analysis is to be financed. The CARES Act was financed simply via government debt, the implications of which remain to be ascertained. In our model, we could easily assume that the government pays for the e_t increases by issuing debt that is to be repaid in a time frame beyond the horizon we are interested in.

5.2 Implications

Figure 5 illustrates the implications of the CARES Act for our model. Upon the imposition of the CARES Act, vulnerable decision makers choose to stay home for an initial period of about a month, after which they start returning to working outside gradually. This process is reversed, however, at about day 130, when outside work participation drops for a couple of weeks. Finally, at about day 150, decision makers return fully outside.

This evolution reflects the interaction between the dynamics of infection, individual decisions, and the financial incentives implicit in the CARES Act. The key aspect of the Act is that the financial reward to outside work relative to staying at home becomes tiny for a while. Then, as soon as the Act is implemented, vulnerable decision makers choose to stay at home. Note that this transition is quite abrupt, which is consistent with the large increase in the unemployment rate in the US in April 2020 which, as Robert Hall has emphasized, reflected an increase of layoffs of workers with jobs rather than job destruction. Also, it is notable that the transition happens even when the rate of infection outside the home is quite small (see upper middle panel).

The fact that decision makers increase slightly their time outside by day 70 is not driven by the incentives of the CARES Act but by the fact that the probability of infection has fallen practically to zero at that moment. Consequently, even though the financial benefit of working outside is close to zero, the likelihood of getting infected is even lower that staying at home 100% is not optimal.

Interestingly, there is a first small wave of infection, that makes it look like the virus will subside by day 100, followed by a second, much bigger wave. Recall that, in the benchmark calibrated model, all infection occurs outside the home since no one stays at home. With the CARES Act, vulnerable decision makers stay at home initially where number of close contacts is smaller, and only essential workers (vulnerables that must work outside) are exposed outside. Hence the first wave is relatively small.

As that wave subsides, people start returning to work outside, responding to the fall in the infection rate. These decisions generate a second infection wave. The rise in that wave induce vulnerable decision makers to again increase their time at home. But then the financial incentives of the CARES Act expire. The consequence is that vulnerables go back to working outside full time. This increases the infection rate, which then crests at about day 190. The number of active cases also increase and peaks at about day 200.

Hence the model generates a case of multiple waves not unlike what has been seen in the US. To underscore this fact, Figure 6 displays the model predictions against US data on cases and deaths. Aside from a time shift, which may reflect that the initial seed in the US was earlier than we are assuming, the dynamics predicted by the model is qualitatively similar to the observed data. Similarly, Figure 7 compares the model predictions for time spent outside the home against mobility data, measured by Apple's driving and walking directions. The data shows that, after an initial drop in 50 percent in mobility, activities outside the home recovered gradually and by August, outside activity was higher than at the beginning of the pandemic. Here also our model implications are at least qualitatively consistent with the evidence.

Figure 8 displays how the incentives embedded in the CARES Act affect the trajectory of the pandemic. The lower left hand panel displays the evolution of the financial incentive to work outside, given by $u'(w_t)(w_t - e_t)$. The CARES Act reduced this incentive to virtually zero for four months. As a consequence, the marginal value of outside work fell during that period, as shown in the upper left panel. When the CARES incentive expire, however, the value of outside work jumps up, prompting vulnerable decision makers to abruptly leave their homes. This, in turn, caused an increase in infection rates, and a second, bigger wave of the pandemic.

It bears emphasizing that the part of the CARES Act that matters for the evolution of the pandemic is only that which changes the relative rewards from working outside the home versus staying home, given by the difference $w_t - e_t$ in our model. In contrast, a policy of transfers to all agents in the population would have essentially no impact on equilibrium dynamics, since it would leave $w_t - e_t$ unchanged. So again, our analysis underscores that fiscal policy can have an impact on the dynamics of a virus, but that impact is determined by the incentive effects embedded in the policy.

6 Application: Social Distancing

Social distancing and mandatory mask wearing are two policies that are used to limit the transmission of the virus between people. In our model, social distancing might be captured by assuming a lower number of close contacts outside the home (lower ρ^m). Mask usage, in turn, can be modeled by a lower probability of transmission per contact with an asymptomatic (lower γ).

In epidemiological models, these policies imply a lower and later peak of active cases and infection rates due to a slower pace of transmission. Moreover, if the these policies are highly effective, it is possible that the arrival of a infectious virus never turns into an epidemic or pandemic.

In our model, however, the impact of this type of policies is ambiguous, due to their impact on individual incentives. By reducing the probability of infection associated with outside activities, social distancing and mask wearing have two opposite effects on behavior. On one hand, by directly reducing current infection risk outside the home, the policies induce agents to spend more time outside. On the other hand, the policies also result in an expected future path with lower probabilities of infection, which increases the future value of being vulnerable relative to being hospitalized, thus raising the expected marginal cost of getting infected today and reducing incentives to work outside.

To illustrate, we analyze the effect of social distancing policies in several alternative cases.

The implications of mandatory masks are qualitatively not different from social distancing. We start by solving the benchmark calibration (log utility) with the effectiveness of social distancing, given by the number of close contacts outside the home, is set at three levels: 5 percent, 25 percent, and 50 percent of the benchmark value of ρ^m. The main outcomes are displayed in Figure 9. In the benchmark calibration, social distancing is not sufficient to induce people to reduce outside time, regardless of effectiveness level. On the other hand, by lowering ρ^m, social distancing directly reduces the pace of transmission. The result is a more favorable virus dynamics, with lower peaks and fewer deaths.

While Figure 9 shows that social distancing does not induce people to reduce their time outside, it does change the incentives that vulnerable decision makers face. This is illustrated in Figure 10. The lower center and right panels of this figure show that greater effectiveness in social distancing implies a lower peak for current net infection risk and, at the same time, a higher expected marginal cost of infection during the initial phase of the pandemic. These two effects go in opposite directions, but with the benchmark calibration, the first one dominates in every case.

Intuitively, therefore, social distancing might induce decision makers to reduce time outside if the expected marginal cost of infection were greater. To check this conjecture, we modify the previous exercise by making the utility cost of death D three times larger than in the benchmark. As discussed before, a greater D increases the incentive to stay home because it reduces the *value* of being hospitalized.

Figure 11 shows our findings. In this "$3D$" scenario, people choose to stay more time at home if social distancing effectiveness is 25 percent. The result is intuitive. If social distancing is less effective, the incentives effect is not strong enough to induce people to reduce outside work time. If social distancing is very effective, infection risk is driven to negligible levels. Relative incentives are displayed in Figure 12.

This analysis underscores that the decision of going outside or staying at home depends on the "double relative" rather than on the absolute values of infection probabilities. It is worth noticing that, in the $3D$ scenario, individual decisionmaking ends up complementing social distancing. That is, social distancing reduces the severity of the pandemic not only by lowering transmission rates, but by inducing individuals to stay at home. In particular, with a 25 percent effectiveness, accumulated deaths, which add up to 1% of the initial population in the benchmark scenario, fall to 0.8% in the $3D$ scenario.

That incentives and individual decisionmaking can complement social distancing depends, however, on the specific scenario under analysis. In some cases, incentive effects can offset social distancing. To illustrate this possibility, we examine a case in which the cost of death is ten times the benchmark. An implication is that, with no social distancing, vulnerables choose to reduce time working outside. Moreover, we assume that the introduction of the social distancing policy is unexpected by agents, and occurs at day 40 after the arrival of the virus.

Figure 13 shows the implications for time working outside, active cases and accumulated deaths. In this case, the implementation of social distancing induces agents to increase time working outside. Moreover, the more effective the policy, the greater is the increase of outside activities.

Notably, despite driving agents to spend more time at the market, social distancing policies still manage to produce less severe epidemiological results. Active cases peaks and accumulated deaths for the three efficiency levels of social distancing are below those in the absence of social distancing. Thus, in this calibration, agents do not increase their outside activities enough to raise the transmission of the virus over the no policy state.

Finally, it is worth mentioning that a basic SIR model would ignore that, after the implementation of a social distancing policy, agents reevaluate their behavior. By doing so, SIR-type models would overestimate the effect of such policy on epidemiological results. This is shown in Figure 14. Given a social distancing policy, the peak of active cases under SIR-type assumptions (dotted lines) would be lower, no matter how effective the policy. This is, of course, because in our model the effect of social distancing is partly offset by a change in individual behavior: people respond to the policy by spending more time outside, which limits the impact of the policy on the number of active cases (dashed lines). In fact, if social distancing has only 5 percent efficiency, the effect of the policy on equilibrium is minimal.

7 A Social Planning Problem

In any equilibrium of our model, vulnerable decision makers choose p_t balancing relative costs and benefits, including those related to the virus and contagion, which individual decision makers take as given. On the other hand, contagion probabilities depend on the distribution of people at home and outside, which is determined by p_t. As a consequence, there are externality effects, and the equilibrium outcome may be socially suboptimal.

To investigate this issue, we consider the case in which, at the beginning of time, the sequence $\{p_t\}$ is chosen to maximize the expected welfare of the typical vulnerable agent (and, therefore, of nearly all the agents in this economy). The problem can be written in recursive form. Let U_t denote period $t's$ social utility, so that:

$$U_t = s_t[qu(w_t) + (1-q)u(p_t w_t + (1-p_t)e_t)] + z_t u(w_t) + [\sum_{i=1}^{H} x_t^{(i)}]u(e_t) - J_t D$$

where, for convenience, we have defined $J_t = \mu x_{t-1}^{(H)}$ as the number of deaths in period t. The first term in the sum in the right hand side is the utility of the s_t vulnerables, which depends on p_t. The other terms gather the utility of z_t recovered and $\sum_{i=1}^{H} x_t^{(i)}$ hospitalized people, minus the cost of J_t deaths.

The value function associated with the planning problem can now be written as $V(s_t, h_t, z_t, J_t, x_t) \equiv V_t$, and the Bellman can be written as

$$V_t = Max_{\Phi_t, p_t \in [0,1]} U_t + \beta V_{t+1}$$

where we take as the period t choice variables the time allocation decision p_t and the probability of infection of a healthy vulnerable, denoted by Φ_t and given by

$$\Phi_t = (q + (1-q)p_t)\phi_t^m + (1-q)(1-p_t)\phi_t^n \tag{15}$$

together with 12, 13, and 14. Meanwhile, state variables evolve according to rewritten 2, 3, 4, 6, and 7 using 15.

This way of writing the planning problem is helpful to sheds some light on the discrepancies between equilibrium outcomes and social optima. In particular, one finds that the social marginal value of p_t is given by:

$$s_t(1-q)[u'(c_t)(w_t - e_t) - \beta\kappa h_t\left(v_{st+1} - v_{ht+1}\right)](\phi_t^m - \phi_t^n) - \Gamma_t$$

where the term Γ_t is given by

$$\Gamma_t = \lambda_t[(q + (1-q)p_t)\frac{\partial \phi_t^m}{\partial p_t}] + (1-q)\left[s_t\beta\kappa h_t\left(\frac{\partial V_{t+1}}{\partial s_{t+1}} - v_{st+1}\right) + \beta\frac{\partial V_{t+1}}{\partial h_{t+1}}\frac{\partial h_{t+1}}{\partial \Phi_t}\right](\phi_t^m - \phi_t^n)$$

with λ_t denoting the Lagrange multiplier associated with 15

Comparing the preceding expressions against the corresponding expression for individual decision makers (equation 10), we see that the term Γ_t captures the externalities involved in the choice of p_t. Individuals ignore that an increase in the time they spend outside, p_t, has a contemporaneous impact on the probability of infection of healthy vulnerables, Φ_t. The social cost of that distortion is given by the first term in the RHS, with the shadow cost of that increase given by λ_t. The second term in the definition of Γ_t expresses the dynamic aspect of the externality. The current choice of p_t has an impact on the evolution of the different population groups, and in particular it affects the number of healthy vulnerables, $s_{t+1}h_{t+1}$.

The discrepancies between the planning solution and the equilibrium outcome in the benchmark case are illustrated in Figure 15.[12]The planning solution differs from the equilibrium

[12]For simplicity, we calibrate Ms equal to M

outcome (which, remember, is also the outcome of the SIR model) in substantial ways. For the first forty days, the planner allows agents to work fully outside, as in the equilibrium model. But then the planner sets p_t to almost zero for about two weeks. After that period, which resembles observed lockdowns, the planner gradually allows vulnerables to return outside. Full return to outside work is not observed until after 260 days since the onset of the pandemic.

Notably, the planning solution reduces outside activities only after 40 days have passed since the outbreak. In terms of the epidemic, it does so after nearly five percent of the population has contracted the virus. More than the particular number, this result suggests that it is not optimal to shutdown outside activity too early.

Likewise, the planning solution underscores that, after an initial strict lockdown phase, a rapid return to outside activity is not optimal. The reason is that the share of the population that remains vulnerable is high and, thus, susceptible to another exponential outbreak of the virus. A smooth return keeps the economy's infection rate on a desired path.

The epidemiological results of the planning solution imply a peak of about 30% of the population in active cases, and a number of deaths of about one half of one percent of the initial population. Hence the planning solution reduces outside activity to control the virus and improve health results. On the other hand, the planner also accepts some health related costs and deaths and avoids a full economic shutdown.

In general, the planning outcome will not be a competitive equilibrium. This means that a government may experience difficulties if it attempts to implement the planning solution, for example by imposing lockdowns. In such a case, vulnerable decision makers would find it individually optimal to deviate from the planning solution. Such tensions seem to be present in practice. In particular, for a sample of 120 countries, Levy-Yeyati and Sartorio (2020) find that lockdown compliance decreases over time, with a stronger trend in economies with higher levels of labor precariousness, that is, economies where most jobs cannot be done from home nor there is a strong safety net.

Hence, in practice, the individual incentives to deviate from the planning solution will vary over time and across countries. This seems to be consistent with our model. For example, the Levy-Yeyati and Sartorio concept of labor precariousness might correspond, in our model, to a greater wedge between w and e, and a greater relative gain from outside work over staying at home. Intuitively, this would make it more difficult for the government to enforce a lockdown. To be able to say more, however, we would need to extend the model to allow for imperfect lockdown enforcement, which is best left for future research.

8 Pandemics and Animal Spirits

Any equilibrium of our model reflects the decisions of individuals in response to the environment they face, including the evolution of the pandemic. But such evolution, as we have seen, depends on whether decision makers stay at home or work in the market. This interaction implies that forward looking behavior and expectations can play a crucial role, and in fact they can lead to multiplicity of equilibria. This is striking, especially when our analysis is compared with that of SIR models, which are purely backward looking. This section discusses how multiple equilibria can emerge, and presents examples. We also argue how policy intervention can help eliminating inferior equilibria, acting as a coordination device.

To identify conditions under which multiple equilibria are likely, suppose that, for a given calibration, there is an equilibrium in which $p_t = 1$ for all t. We will refer to that equilibrium as being of the *SIR-type* since, as seen in subsection 4.2, it replicates the outcome of a naive SIR model.

For the given calibration, we can ask: is there a different equilibrium? If the answer is yes, it must be that $p_t < 1$ for at least one t. In other words, there must be at least one period in which vulnerable decision makers reduce their time outside their homes, presumably because of fear of infection. For concreteness, we will say that this equilibrium is of the *precautionary type.*

Now, since the SIR-type equilibrium is assumed to exist, (10) implies that, for each t,

$$u'(w_t)(w_t - e_t) \geq \beta\kappa\hat{h}_t(\hat{\phi}_t^m - \hat{\phi}_t^n)(\hat{v}_{st+1} - \hat{v}_{ht+1})$$

where we are using carets to identify endogenous variables in the SIR-type equilibrium (i.e. $\hat{p}_t = 1$, all t). In words, this expression says that, in every period, the current utility net gain of spending more time outside home is equal or greater than the expected marginal cost of infection risk, even if the individual is already fully outside the home.

Conversely, in a precautionary equilibrium, which we identify with tildes, for any t in which $\tilde{p}_t < 1$ it must be true that:

$$u'(w_t)(w_t - e_t) < \beta\kappa\tilde{h}_t(\tilde{\phi}_t^m - \tilde{\phi}_t^n)(\tilde{v}_{st+1} - \tilde{v}_{ht+1})$$

If a precautionary equilibrium exists alongside the SIR-type equilibrium, let τ denote the first period t in which $\tilde{p}_t < 1$. Since we assuming a fixed calibration, for every period below and equal to τ the population shares of epidemiological states will be the same in both equilibria. More formally, $\{\tilde{h}_t, \tilde{s}_t, \tilde{x}_t, \tilde{z}_t, \tilde{\omega}_t\}$ is equal to $\{\hat{h}_t, \hat{s}_t, \hat{x}_t, \hat{z}_t, \hat{\omega}_t\}$ for all $t \leq \tau$. This result includes period τ because shares of epidemiological states are determined by decisions in the previous period. Combine this observation with the two previous inequalities in period $t = \tau$ to obtain

$$(\tilde{\phi}_\tau^m - \tilde{\phi}_\tau^n)(\tilde{v}_{s\tau+1} - \tilde{v}_{h\tau+1}) > (\hat{\phi}_\tau^m - \hat{\phi}_\tau^n)(\hat{v}_{s\tau+1} - \hat{v}_{h\tau+1})$$

To see what this condition implies, recall that, for a calibration with a constant sequence of home rewards ($e_t = e$, all t), v_{ht} is a constant v_h, independent of which equilibrium obtains. Also, because $\hat{h}_\tau = \tilde{h}_\tau$, $\hat{\phi}_\tau^n = \tilde{\phi}_\tau^n$, while $\hat{\phi}_\tau^m \geq \tilde{\phi}_t^m$, since ϕ_t^m is increasing in p_t. All of these facts together imply that the preceding inequality can hold only if the value of remaining vulnerable is sufficiently greater in the precautionary equilibrium than in the SIR type equilibrium:

$$\tilde{v}_{s\tau+1} >> \hat{v}_{s\tau+1}$$

This condition underscores that multiple equilibria can emerge, but only in the presence of strategic complementarities that have a sufficiently strong impact on expectations. Suppose that a SIR-type equilibrium coexists with a precautionary equilibrium, and consider the dilemma faced by a vulnerable decision maker at $t = \tau$. In the precautionary equilibrium, she knows that $\tilde{p}_t < 1 = \hat{p}_t$. This actually implies that the risk of infection outside the home is *less* in the precautionary equilibrium than in the SIR-type one (since $\hat{\phi}_\tau^m \geq \tilde{\phi}_t^m$). That the agent reduces her time working outside then reveals that, in her assessment, the expected value of remaining in the vulnerable group in $\tau + 1$ is much larger in the precautionary equilibrium than in the one of the SIR type.

In turn, how much greater the precautionary equilibrium value of starting $\tau + 1$ in the vulnerable group relative to the value in the SIR equilibrium depends on how much future infection probabilities fall when agents spend less time outside their homes from τ on, which determine the expectation of financial loss and death due to the virus.

In sum, the interaction between the trajectory of infection and forward-looking decision making results in the possibility of multiple equilibria. Strategic complementarities may exist, but in our model they must be dynamic ones Expectations can therefore play a key role.

To illustrate, we provide an example of multiple equilibria. We run the benchmark calibration, except that we increase the utility cost of death (D) fivefold. Recall from subsection xx that the benchmark calibration has a SIR-type equilibrium, whereas the scenario with 10 times D has a precautionary-type equilibrium. In both cases, in our numerical work we have been unable to find other equilibria. In contrast, in the scenario with five times D we have found two equilibria, one SIR-type and the other precautionary-type.

The different outcomes in the number of equilibria due to changes in D reinforce the argument that strategic complementarities occur through expectations. Moreover, it also signals that multiple equilibria occur when the expected cost of infection risk, captured by $v_{st+1} - v_{ht+1}$, is neither too big nor too small. If too big, the SIR type equilibrium disappeards because it is too costly to allot your whole time to market activities. If too small, the precautionary equilibrium is not feasible because it is not optimal to sacrifice market time.

Figure 16 shows the simulation results for both equilibria. In this case, τ is equal to 44; also, market participation is full in the precautionary equilibrium after 88 days. Thus, the precautionary equilibrium deviates from the SIR-type equilibrium for over a month. As a result, the precautionary equilibrium reaches a steady state with about one third fewer accumulated deaths than in the SIR-type equilibrium, and with about 20 percent of the population never being infected.

Notably, the infection rate of the economy (center panel, upper row in Figure 16) is the same between equilibria up to day τ. However, at this point, by reducing their time at the market, the infection rate in the precautionary equilibrium stagnates and smoothly falls to zero by day 120.

The path of the economy's infection rate underscores that decision makers are forward looking. Up to day τ, there is no difference in the evolution of the pandemic between equilibria. At that point, however, paths diverge, reflecting differences in expectations. In the precautionary equilibrium, decision makers expect relatively low infection rates in the future, so they find optimal to reduce their market activities. In contrast, in the SIR-type equilibrium, expectations are that people will not reduce their time at the markets, and, thus, infection rates will relatively high in the future. At this point, it is not optimal for a decision maker to reduce their market time since the expected marginal cost of infection risk is low relative to the current market gain. In this sense, emergence of the SIR-type equilibrium is a coordination failure.

Figure 17 examines the incentives of a vulnerable decision maker in both equilibria. The lower row of the figure displays the three parts of the derivative of the agent's maximand. Note that, at day τ (first dotted vertical line), both the net current gain (left panel) and the net infection risk (center panel) are practically the same between equilibria. Thus, whether decision makers decide to reduce their time at the market can be entirely explained by different expectations (right panel). That difference reflects that the expected path of infection rates is lower in the precautionary equilibrium than in the SIR-type equilibrium.

The possibility of multiple equilibria underscores how forward looking behavior and different expectations can potentially have a dramatic impact on the dynamics of a pandemic. This result is not just a novelty of the model, but it also has important consequences for policy.

To be concrete, suppose that a SIR-type equilibrium coexists with a precautionary equilibrium in which time outside the home is denoted by $\tilde{p}_t$. Then a lockdown policy of restricting non-essential work outside the home to be no more than $\tilde{p}_t$ would eliminate the SIR-type equilibrium, while leaving the precautionary equilibrium intact (this is clearly the case, since the restriction would not be binding for vulnerable decision makers in the precautionary equilibrium). Notably, such a lockdown policy would not only eliminate the less desirable equilibrium, but also would not create incentives for individuals to deviate from $\tilde{p}_t$. The lockdown would act as a coordination device, and would not suffer from the same implementation challenges affecting the optimal planning solution, as mentioned at the end of the previous section.

Notably, the same outcome can be implemented not with a lockdown but with economic incentives. One such policy would be to require nonessential workers to pay a marginal tax at rate $(w_t - e_t)$ on earnings from time working outside over and above $\tilde{p}_t$. Intuitively, such a tax would eliminate the incentives for vulnerable decision makers to spend more than $\tilde{p}_t$ time outside their homes. Clearly, the SIR-type equilibrium would disappear, and the precautionary equilibrium would remain – in fact, in the precautionary equilibrium, no one would pay the tax.

We close this section by noting that, even if policy can eliminate the bad SIR-type equilibrium, the surviving precautionary would still be inferior to the social planning outcome of the previous section. In this sense, the ability of a government to use policy as a coordination device is no substitute for the institutional strength needed to implement the social optimum.

9 Final Remarks

In this paper we have explored a very simple idea: how fast and widely a virus spreads across individuals depends on how they behave in response to the health costs of contagion, but also to economic and financial incentives. This idea is simple but quite powerful. It has implications for the dynamics of infection, the evolution of economic activity during a pandemic, the role of forward-looking behavior and expectations and, perhaps most importantly, the scope for public health directives and economic policy to influence both lives and livelihoods during a pandemic.

The model here can be extended in multiple interesting directions. We comment on the ones that appear most interesting or important to us.

Today the rollout of vaccines is the best hope to put an end to the pandemic. Indeed, the quick development of vaccines has been a major triumphs for science. Yet it is plausible (at least, after having read this paper) that the expectation of vaccine development may have played a role in extending the duration of the Covid-19 episode.

Imagine a situation in which a vaccine can appear in each period with some given probability. Once the vaccine appears, a fraction of the vulnerable population in each remaining period receives a jab and it becomes immune, joining for all practical purposes the recovered population. In our model, vaccine development and distribution efforts could affect individual incentives to work outside or stay at home, and may generate effects in opposite directions. For instance, the imminent arrival of a vaccine could lower contagion probabilities and induce more risk-taking behavior. But, at the same time, it could prompt expectations of a vigorous economic recovery, increasing the value of remaining healthy and reducing risk-taking behavior. We believe that this is a promising question for future research.[13]

We have followed convention in economics and assumed that agents have perfect foresight. But that assumption, never innocuous, becomes even more questionable during a pandemic, when individuals and households are subjected to new shocks and experiences for which there is

[13] The importance of this issue can be illustrated by the current experience in Chile, whose vaccination drive has been among the most successful in the world, in spite of which infection rates have accelerated. Observers agree the explanation is that, reassured by the vaccination success, Chileans returned to outside activities.

little precedent and limited accumulated knowledge. At the beginning of the Covid-19 episode there was considerable uncertainty about the basic parameters regarding virus contagion, for example, and much debate on the effectiveness of masks and other prophylactic devices. One might then ask how much our analysis might change if one departs from the perfect foresight assumption and allows for a learning process to take place.

In building the model of this paper we have made every effort to rely on first principles. One payoff is that the model can be modified in straightforward ways to tackle many important issues. A possible extension could examine the implications of heterogeneity over the $w_t - e_t$ wedge. This may capture inequalities in the ability of different people to work from home. Alternatively, a distribution over w_t might capture income inequality. The model might then provide some lessons on how inequalities in income, wealth, or labor market opportunities affect the severity of a pandemic. In that way, the model would not only give insights about what policies are most effective to control a pandemic once it has started, but also could help identify the socio-economic characteristics and institutions that left countries better or worse prepared to face the Covid 19 pandemic –and which will be highly relevant when the next threatening virus comes along.

References

[1] Acemoglu, D., Chernozhukov, V., Werning, I., & Whinston, M. D. (2020). A multi-risk SIR model with optimally targeted lockdown (No. w27102). National Bureau of Economic Research.

[2] Alvarez, F. E., Argente, D., & Lippi, F. (2020). A simple planning problem for covid-19 lockdown (No. w26981). National Bureau of Economic Research.

[3] Baldwin, R. and B. Weder di Mauro, editors, Economics in the Time of COVID-19, CEPR Press, 2020.

[4] Chang, R., & Velasco, A. (2020). Economic policy incentives to preserve lives and livelihoods NBER Working Paper 27020

[5] Chen, F. (2012), A Mathematical Analysis of Public Avoidance Behavior During Epidemics Using Game Theory, Journal of Theoretical Biology, 302, 18-28.

[6] Eichenbaum, M. S., Rebelo, S., & Trabandt, M. (2020). The macroeconomics of epidemics. NBER Working Paper 26882

[7] Faria e Castro, M., "Fiscal Policy During a Pandemic". Working Paper 2020-006, Federal Reserve Bank of Saint Louis, March 2020.

[8] Fornaro, L. and M. Wolf, "Covid-19 Coronavirus and Macroeconomic Policy". Technical Report, CEPR, March 2020.

[9] Ganong, P., Noel, P. J., & Vavra, J. S. (2020). US Unemployment Insurance Replacement Rates During the Pandemic (No. w27216). National Bureau of Economic Research.

[10] He, D., Zhao, S., Lin, Q., Zhuang, Z., Cao, P., Wang, M. H., & Yang, L. (2020). The relative transmissibility of asymptomatic cases among close contacts. International Journal of Infectious Diseases.

[11] Jones, C., T. Philippon and V. Venkateswaran, "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home". Manuscript, NYU, April 6, 2020.

[12] Jorda, O., S. Singh, and A. Taylor, "Longer Run Economic Consequences of Pandemics". Working paper 2020-09, Federal Reserve Bank of San Francisco, 2020.

[13] Kermack, W. O. and McKendrick, A. G. "A Contribution to the Mathematical Theory of Epidemics". Proceedings of the Royal Society A. 115 (772), 1927.

[14] Kniesner, T. J., & Viscusi, W. K. (2019). The value of a statistical life. Forthcoming, Oxford Research Encyclopedia of Economics and Finance, 19-15.

[15] Loayza, N. and S. Pennings, "Macroeconomic Policy in the Time of COVID-19: A Primer for Developing Countries". Research and Policy Briefs Number 28, World Bank Malaysia Hub, March 2020.

[16] Mossong J, Hens N, Jit M, Beutels P, Auranen K, Mikolajczyk R, et al. (2008) Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases. PLoS Med 5(3): e74. https://doi.org/10.1371/journal.pmed.0050074

[17] Verity, R., Okell, L. C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N., ... & Dighe, A. (2020). Estimates of the severity of coronavirus disease 2019: a model-based analysis. The Lancet infectious diseases.

[18] Weiss, H. H. (2013). "The SIR model and the foundations of public health". Materials Matematics, 0001-17 Volum 2013, treball no. 3, 17 pp. ISSN: 1887-1097

[19] Yeyati, E. L., & Sartorio, L. (2020). Take me out: De facto limits on strict lockdowns in developing countries. Covid Economics, 59.

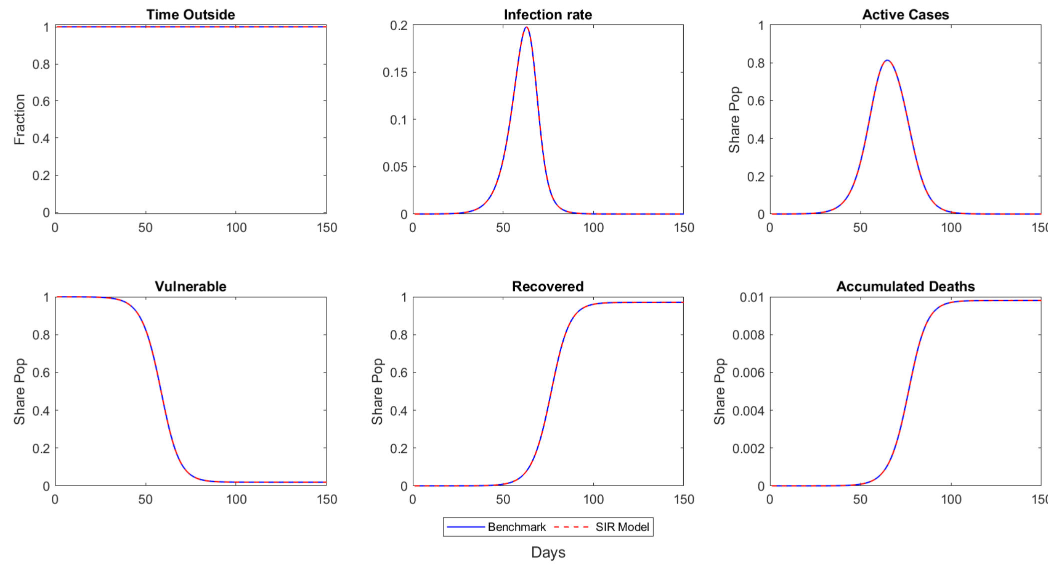

Figure 1: Model (log utility) versus SIR

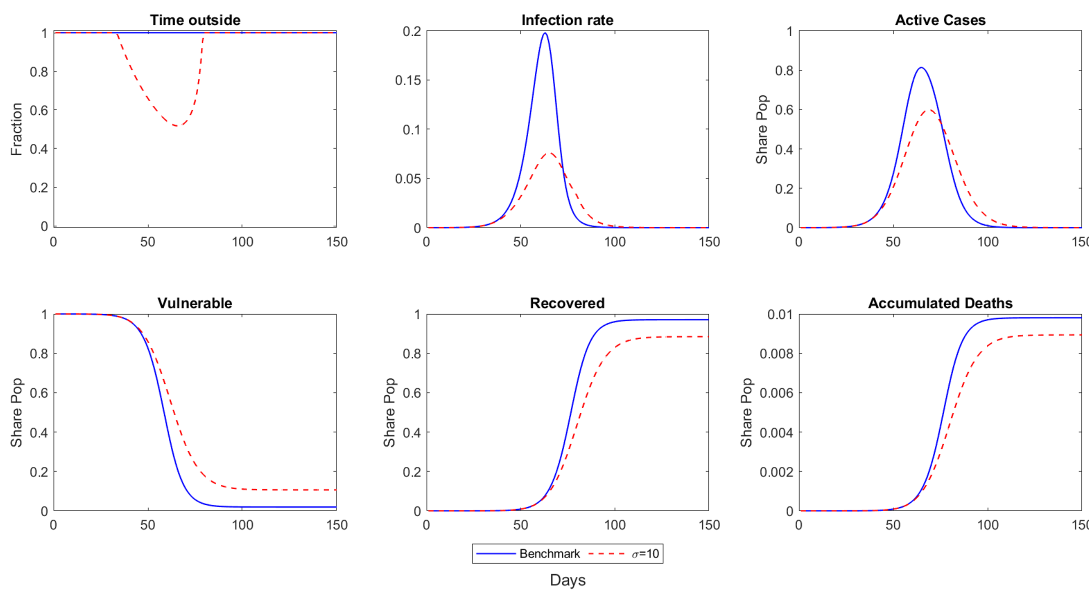

Figure 2: Model ($\sigma = 10$) vs SIR

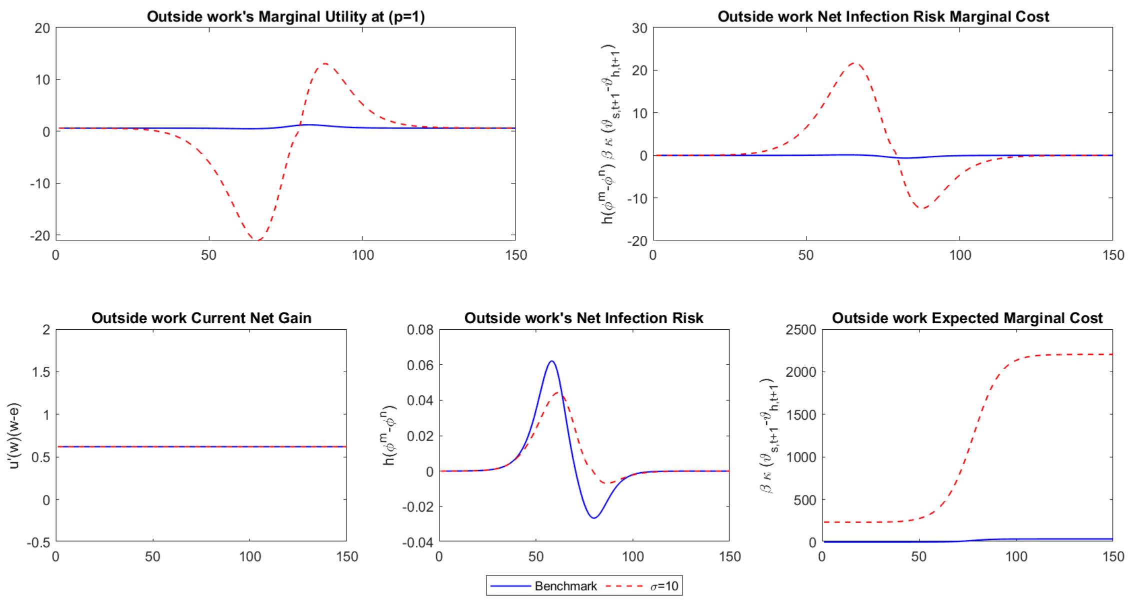

Figure 3: Incentives for Decision Making ($\sigma = 10$)

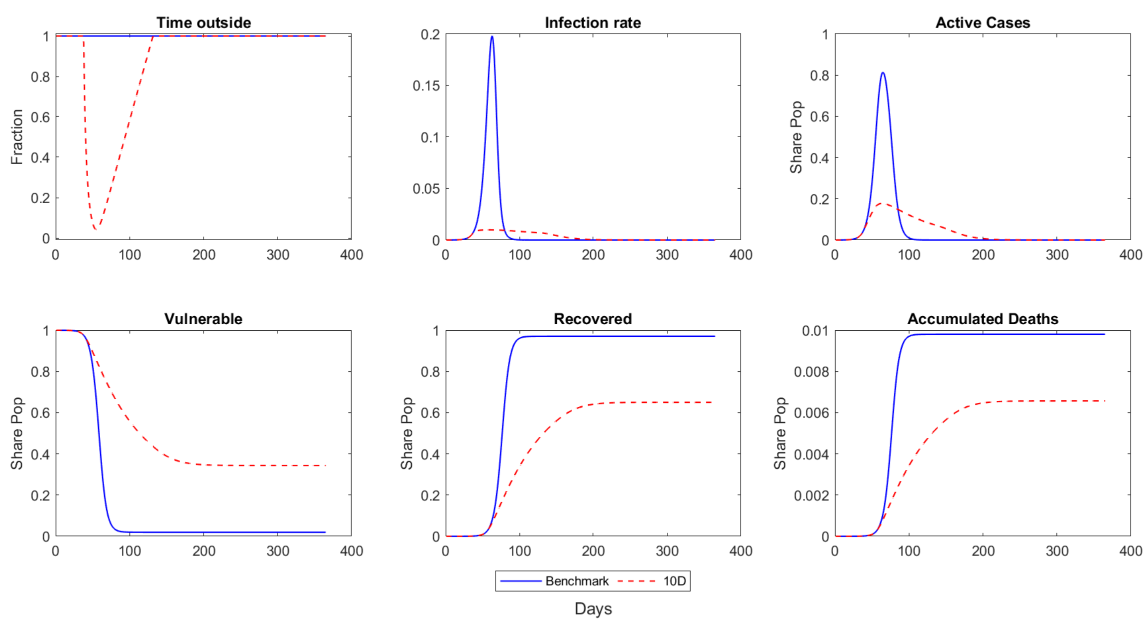

Figure 4: Bigger Fear of Death

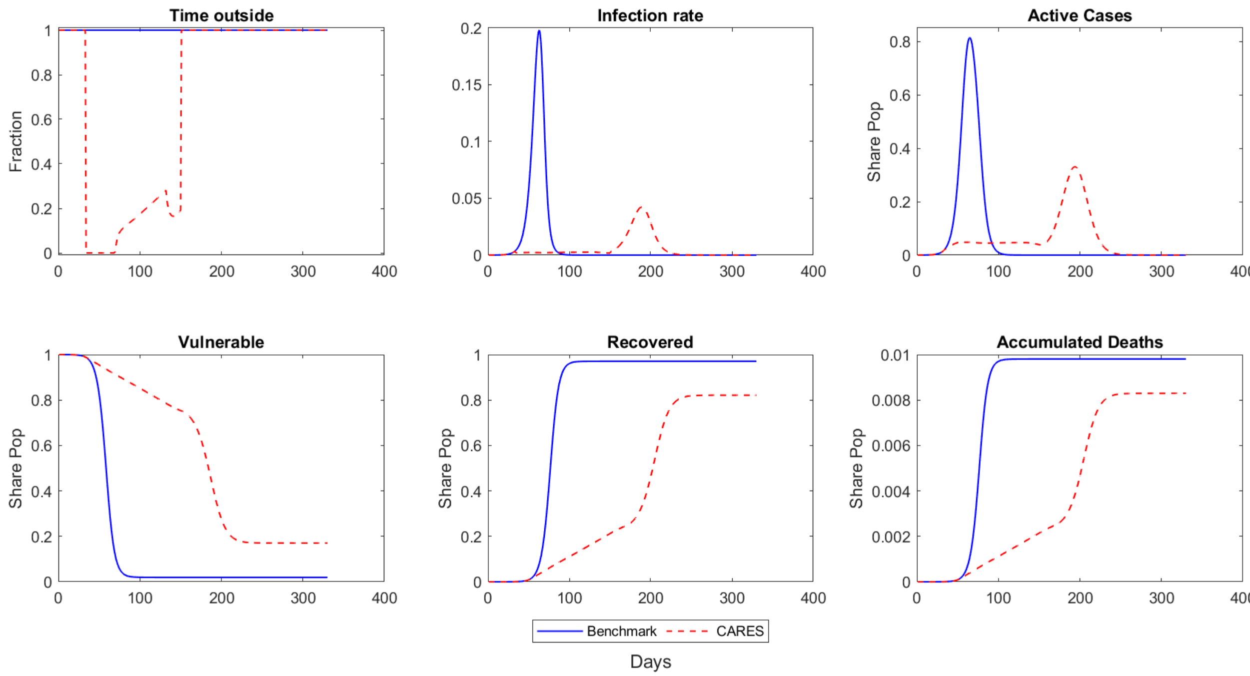

Figure 5: CARES and Multiple Waves

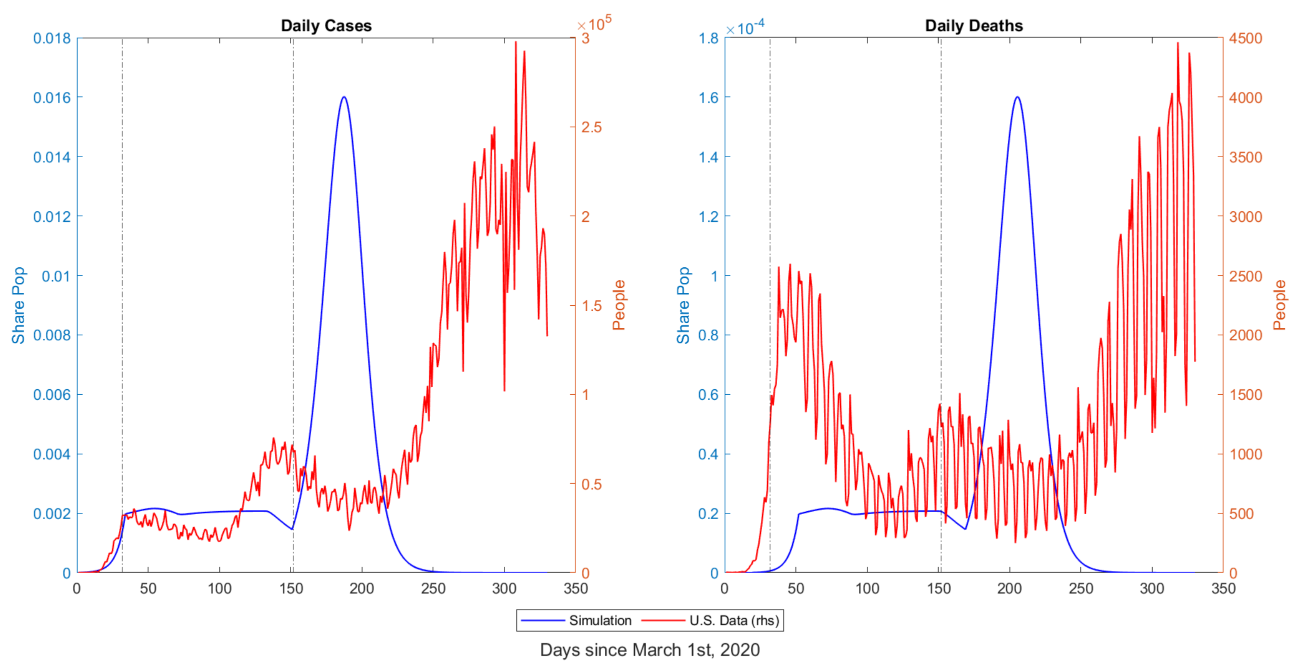

Figure 6: Cases and Deaths: Model and US data

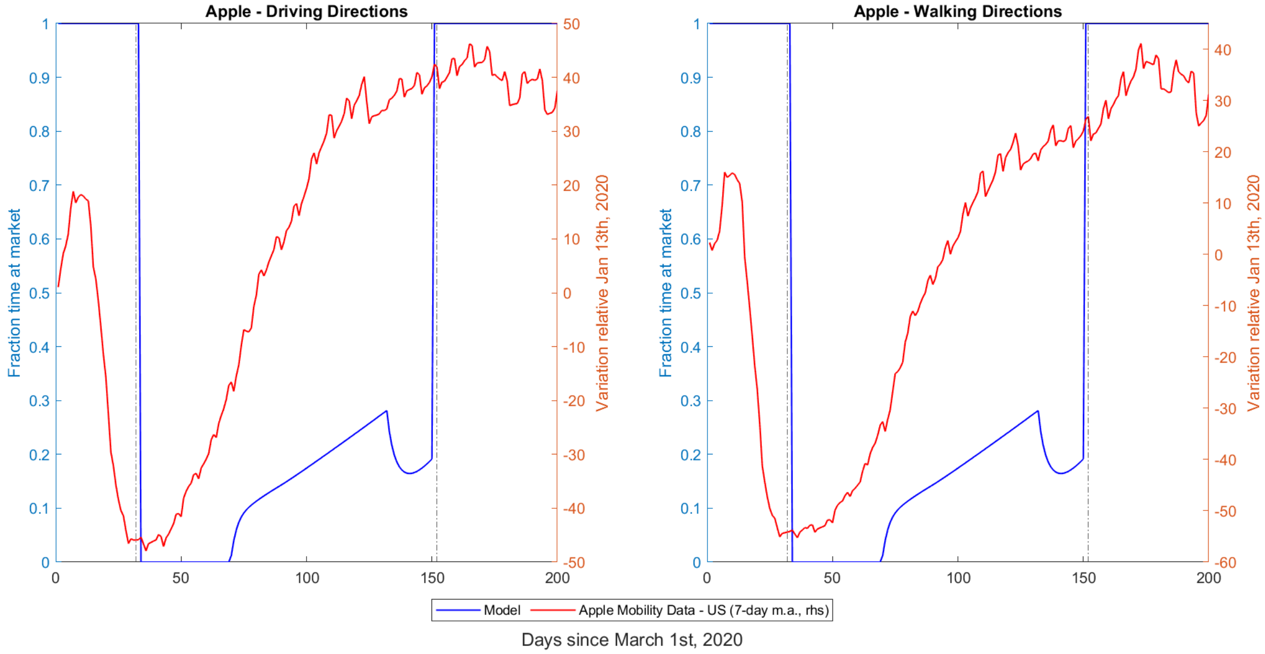

Figure 7: CARES and Mobility: Model and US Data

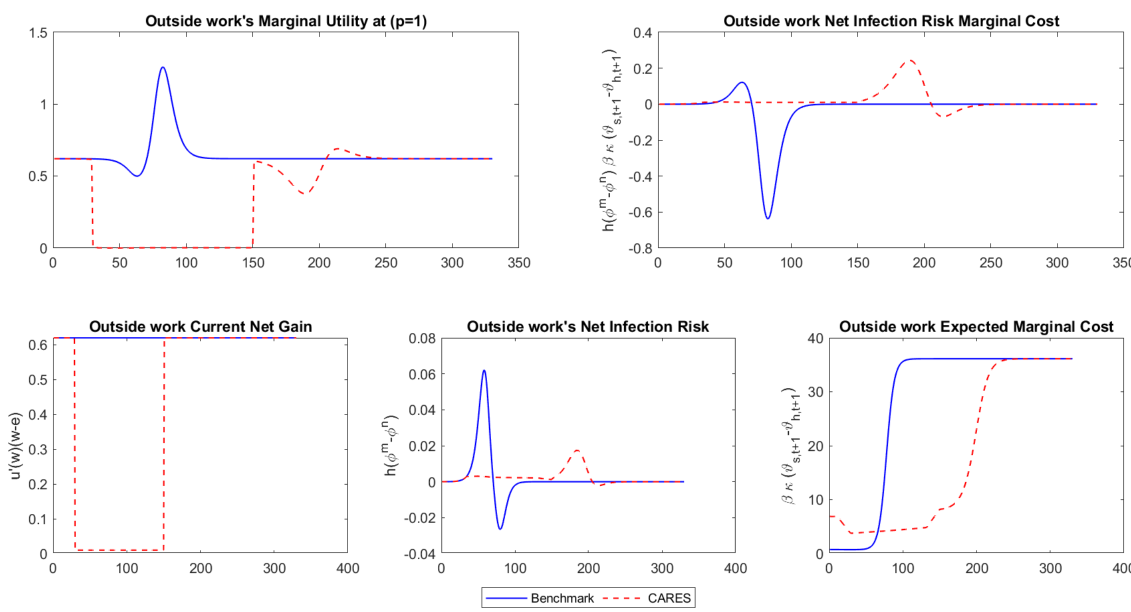

Figure 8: Incentives in the CARES Act

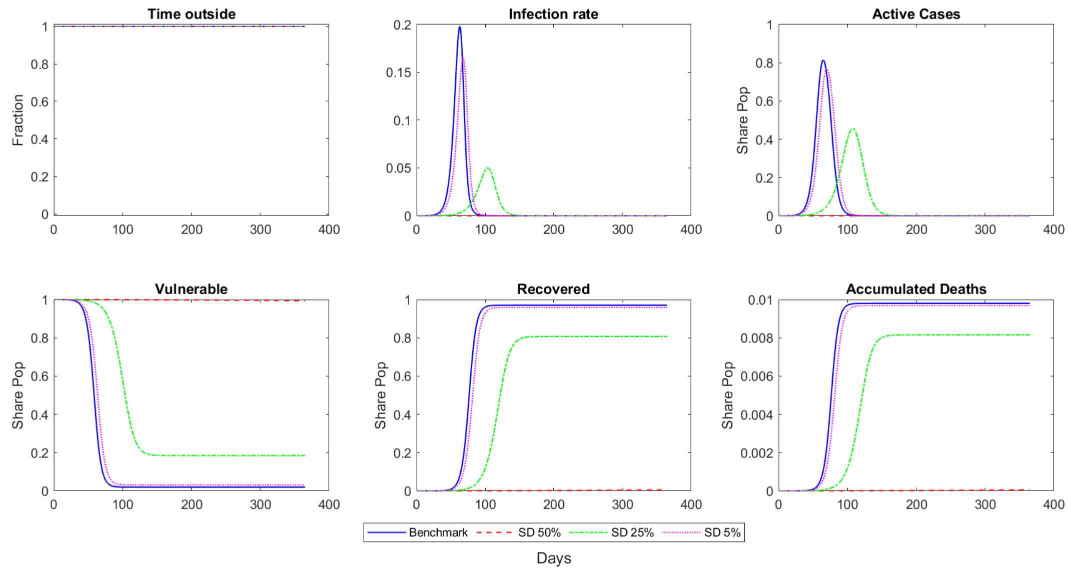

Figure 9: Social Distancing

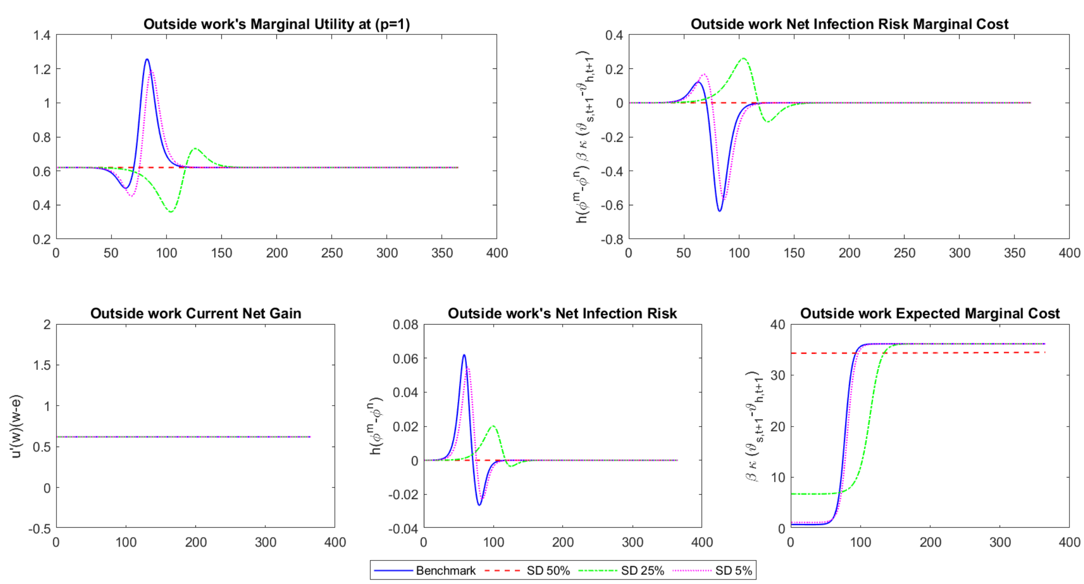

Figure 10: Incentives and Social Distancing

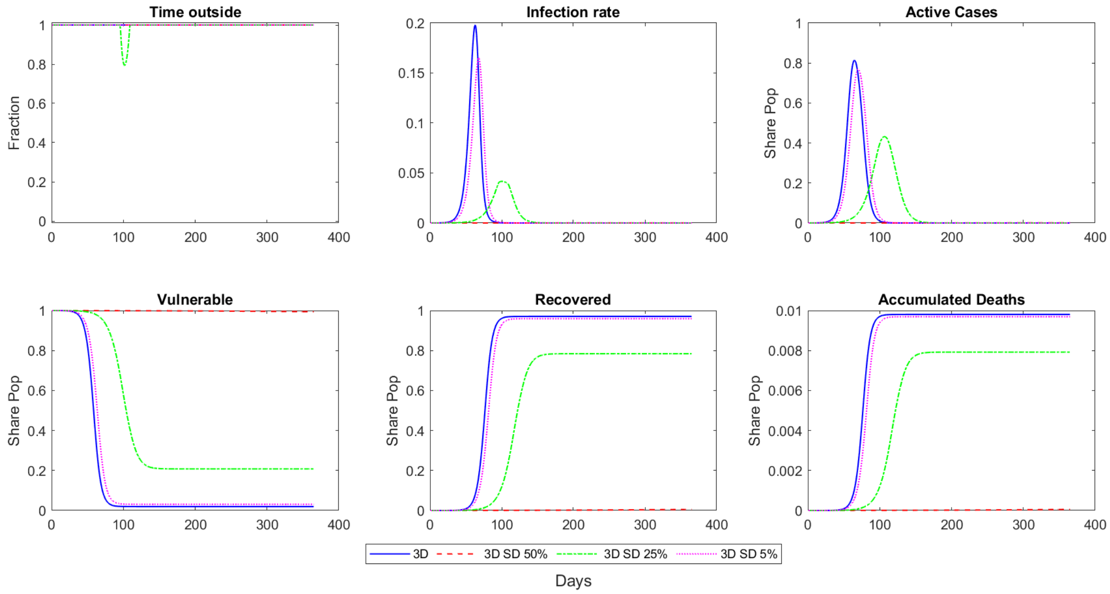

Figure 11: Social Distancing With Cost of Death = $3D$

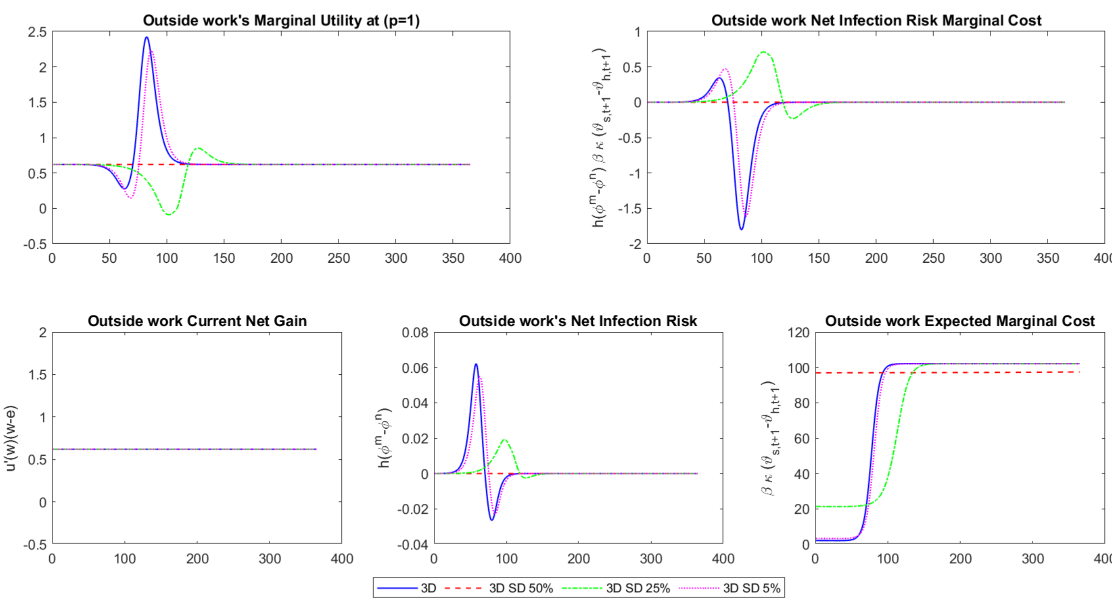

Figure 12: Incentives and Social Distancing ($3D$ case)

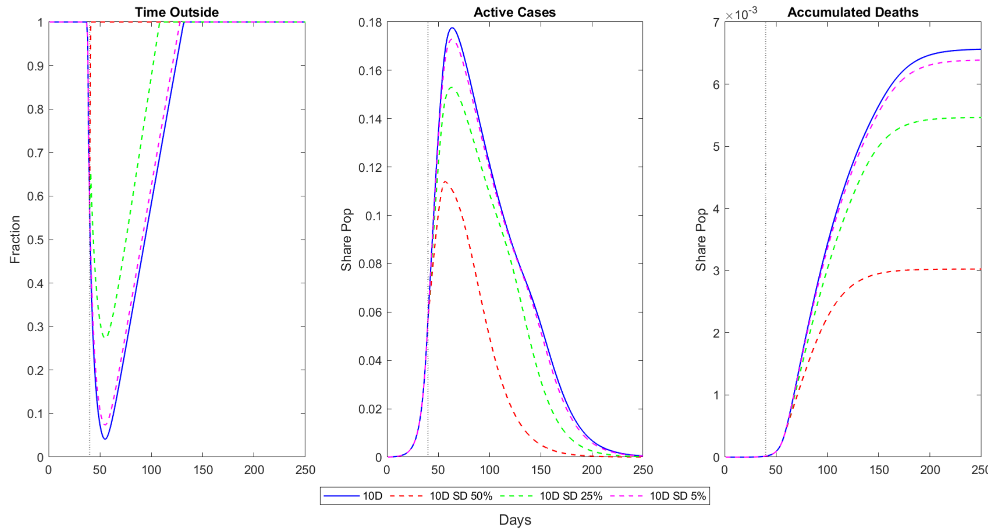

Figure 13: Social Distancing (cost of death = $10D$)

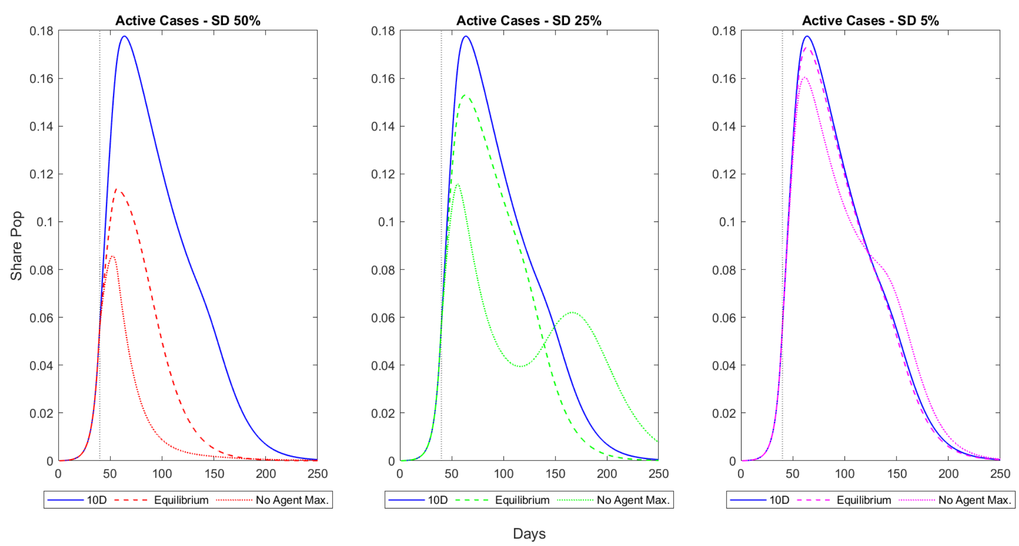

Figure 14: Social Distancing vs SIR

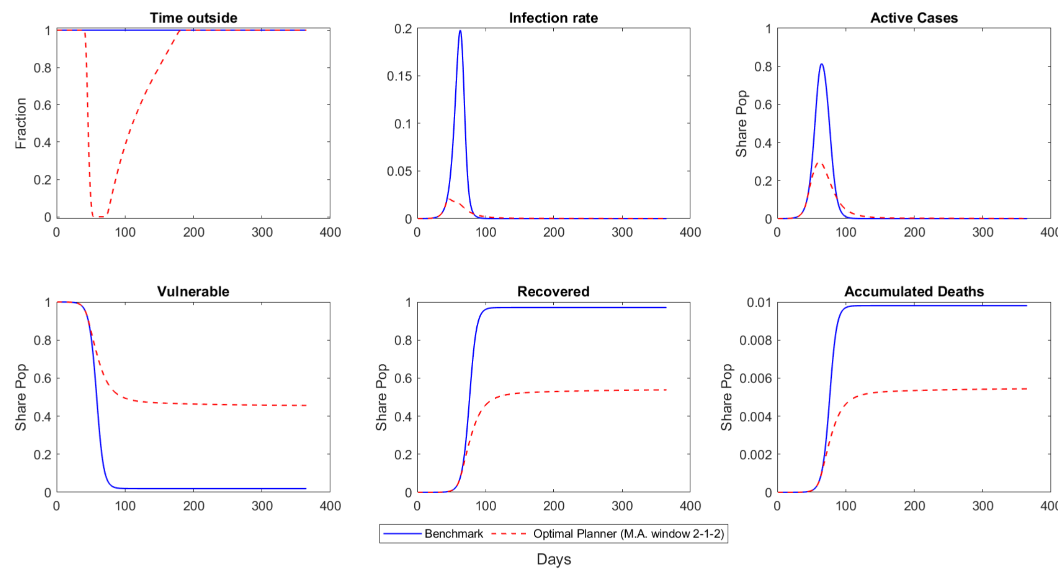

Figure 15: Equilibrium vs Socially Optimal Plan

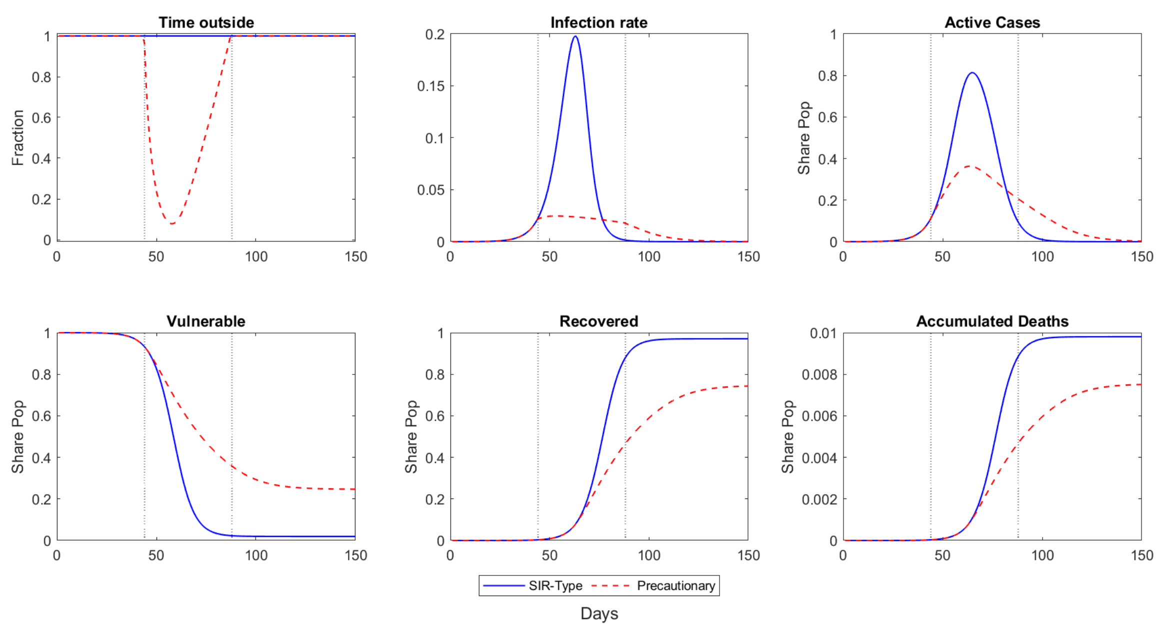

Figure 16: Multiple Equilbria: Pandemic Dynamics

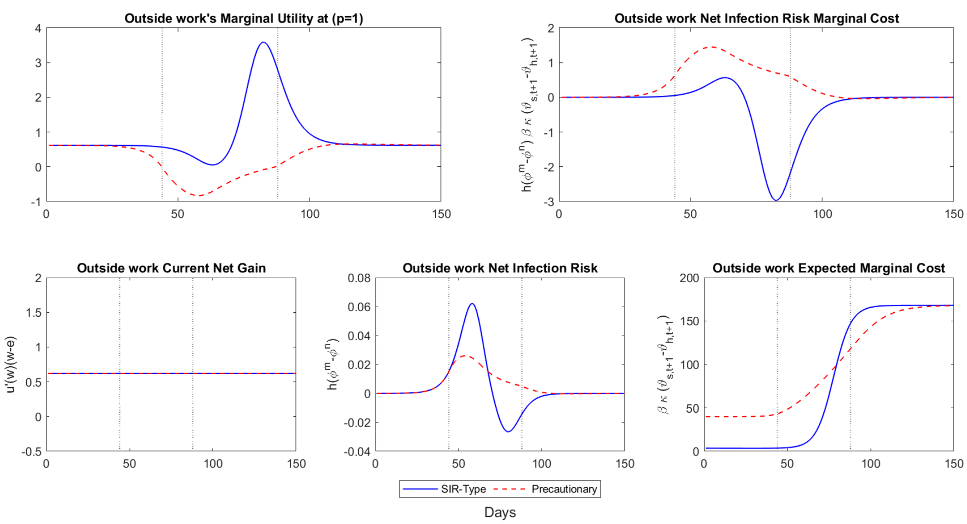

Figure 17: Multiple Equilibria: Individual Incentives