Pandemics, Incentives, and Economic Policy: A Dynamic Model*

Roberto Chang† Humberto Martínez‡ Andrés Velasco§

February 2021

Abstract

The dynamics of a pandemic depend on individual decisions about, for instance, social distancing. In making these decisions, individuals respond to the incentives they face, including economic ones. This implies that economic policy can help not only by alleviating economic losses due to the pandemic but also by influencing the pandemic's trajectory itself via incentives. To investigate this idea, we develop a dynamic equilibrium model of an economy subject to a pandemic.In the model, individuals choose whether to work in the market or stay at home; market participation yields higher current pay but also a higher risk of infection. In turn, infection rates depend on the extent of market participation, as in typical SIR models. We use the model to investigate how pandemic dynamics depend on aspects of the model, such as preferences, which have no bearing in standard SIR models. Most novel, fiscal policy can impact the evolution of the pandemic.To illustrate, we show that the incentives embedded in a fiscal package resembling the U.S. CARES Act can result in two waves of infection.

*Preliminary and Incomplete. We will greatly appreciate any comments and suggestions. This project started while Chang was serving as BP Centennial Professor at LSE, whose hospitality he acknowledges with thanks.

†Rutgers University and NBER. chang@econ.rutgers.edu

‡Rutgers University

§London School of Economics and Political Science

1 Introduction

The Covid-19 pandemic has stimulated a myriad of new research contributions from economists, seeking to understand its economic, social, and political implications, as well as to devise appropriate policy responses to mitigate the damage, especially on the most vulnerable sectors of the population, and to prepare the groundwork for an eventual recovery.

Notably, much of the recent literature treats the pandemic as an exogenous shock, and then investigates policy proposals aiming at alleviating its negative consequences. Doing so, however, ignores the fact that the propagation of the pandemic is, at least to some extent, determined by economic decisions, and hence the size and dynamics of the Covid-19 "shock" may be affected by economic policies themselves. Such omission may have significant implications, in particular, for policy analysis, as emphasized by Chang and Velasco (2020, CV henceforth).

The premise that economic decisions and policies can influence the severity of a pandemic is not a hard one to justify. In particular, the efficacy of health policy directives such as lockdowns, stay-at-home guidelines, and social distancing rules, depends on the decisions of individuals in view of the perceived costs and benefits of compliance. CV developed a simple model in which economic policies affect the cost-benefit analysis and, as a consequence, have an influence on individual decisions and the transmission of a virus.

Following CV, this paper develops and analyzes the implications of a dynamic model of the interaction of virus transmission, individual decisions, and economic incentives and policy. In contrast with CV's model, which assumed two periods, our model here assumes an infinite time horizon, and hence represents an improvement in terms of empirical realism. Also, the assumption of an infinite horizon makes our model more easily comparable to the dominant Susceptibles-Infected-Recovered (SIR) model of pandemics.

As in CV, we focus on a population of individuals that, normally, spend their time working in the market. 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, with luck, they recover, or die if unlucky.

Individuals not showing symptoms of the virus can work in the market or stay at home. Staying at home is costly in terms of lost income, but the probability of infection is, most of the time, greater in the market than at home. This means that agents decide between market work and home on the basis of a "double relative": market income versus home income; and the value of remaining asymptomatic relative to the cost of sickness.

The probabilities of infection in the market and at home, in turn, depend naturally on the numbers of agents in each location and, therefore, on individual decisions. The outcome of the model is, thus, given by a general equilibrium in which location decisions and infection probabilities are simultaneously determined.

Some analytical results are available from inspection of the model equations, as in basic SIR models. The model is highly nonlinear and time dependent, however, so that the resulting dynamics are generally complex and depend on model parameters. To derive additional information, we analyze equilibria with parameters calibrated to match U.S. data. The resulting equilibrium pattern of infections displays the hump shaped behavior characteristic of a pandemic, which is the hallmark of SIR models.

On the other hand, and in contrast with SIR models, the outcomes of our model reflect individual decisions and the corresponding incentives. This means that changes in, say, the "technology" of infection affect equilibrium outcomes not only directly but also due to the induced changes in individual behavior. For example, an increase in the infectiousness of the virus has a direct influence on the speed of transmission, but also means that, to avoid contagion, individuals will be more inclined to stay at home, which acts as an offsetting force.

More consequentially for our purposes, changes in economic variables, which would not have an impact on SIR models, can have an impact on the decisions of individuals and, through them, on the dynamics of virus transmission of our model. This is the case, in particular, of the relative rewards of working in the market versus staying at home. Likewise, the outcomes of our model depend on aspects of individual behavior, such as risk aversion, that have no bearing on SIR outcomes.

In order to emphasize the importance of the interaction between economic incentives and pandemic dynamics, we use the model to examine the implications of a fiscal policy package similar to the U.S. 2020 CARES Act. We feed into the model two essential aspects of the act: an increase in unemployment benefits which virtually eliminated the difference between market income and home income; but did so for a limited time period. In our model, such a policy implies that the pandemic comes in two waves, consistent with the observed evolution of Covid-19 in the US.

In the context of our model, the intuition is clear. While they lasted, the unemployment benefits enacted by the CARES Act changed incentives for people to refrain from working outside the home. This change in behavior was instrumental in helping controlling the epidemic, contributing with a reduction in infections after an initial peak. However, after the CARES Act lapsed, the associated increase in market income relative to home income induced individuals to return to the market, starting a second wave of the pandemic.

Remarkably, in the absence of the CARES Act, our model predicts that the U.S. would have experienced a single peaked pandemic. This would also be the case if individual locations were exogenous, in whose case our model becomes one of the SIR type.

In a separate exercise, we analyze how a social distancing policy affects economic incentives. Results show that this policy generates two opposite effects, and, thus, it has an ambiguous impact on individual behavior. In our model, this is captured by reducing the number of close contacts when individuals decide to go to the market which declines current and future infection probabilities. Unsurprisingly, lower current net infection risk increases incentives for market activities. However, expected infection risk is lower as well increasing the cost of getting infected today and missing out on lower infection rates in the future. This shifts incentives towards staying at home. We illustrate the ambiguous impact of social distancing by providing different calibrations where the equilibrium result of this policy can be an increase of market activities in one while a reduction in another scenario. This exercise underscores that what matters for individual maximization are not absolute values of infection rates but what CV denoted as the double relative.

Our model is built on first principles, so it is especially useful for welfare analysis. To illustrate, we analyze the obvious social planning problem in our model, one in which the social objective is given by the discounted expected welfare of the economy's agents, weighted by their population size. It becomes clear that the model incorporates an externality, as individuals base their choices on their incomes and their own probabilities of contagion, but fail to take into account the impact of those choices on the evolution of the population and, through it, the economy wide dynamics of contagion.

With the benchmark calibration, the planning solution restricts market participation drastically at first, and then allows agents to return to the market gradually. Hence the solution looks like a "lockdown" and contrasts with the decentralized equilibrium, in which market participation is full. The optimal policy reduces infection but does not eliminate it completely. In particular, the number of deaths falls to about one half of one percent of the population.

Beyond comparing optimal versus decentralized outcomes, we ask: what is an individual's best response if the planner were to implement the optimal lockdown? Our results show that, if an individual believes that others will comply with the lockdown, then market infection risk falls and incentives to stay at home decline.enough so that he will find it in his own interest to return to the market fully. The optimal policy may then be politically difficult to implement. Without enforcement or an economic policy that reduces the wedge between market and home income, lockdown compliance can be expected to be, at best, short-lived.[1]

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 the model mechanics. The implications of a fiscal package similar to the CARES Act are discussed in section 5. Social distancing is discussed

[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

in section 6. A social planning problem is proposed in section 7, and its solution is compared with the equilibrium of the model. The relation of our work with the literature and some final thoughts are gathered in section 8.

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*, *market*, and *hospital*. As in CV, 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 in the market in period t produces and receives an amount w_t. We can think of w_t as a wage, 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.[2] 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 work every day in the market, and receive utility:

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

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

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

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 CV. An exogenous random fraction $q > 0$ of the vulnerables is selected and must go to the market. 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 to spend in the market. In the parlance of CV, 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 in the market (resp. at home). Hence a healthy vulnerable that spends a fraction p_t of her day in the market 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 $D \geq 1$ 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 market wages, defined above as v_z. On the other hand, we assume that death involves a utility cost $M \geq 0$.

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

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

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

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

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

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

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

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

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, as recognized in CV, 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 to the market. 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 market 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 in the market 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 between the market or home. 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 in the market, 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 in the market, 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 market 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 going to the market 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, as in CV, decisions to work or stay at home depend on a "double relative": the current payoff to 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

As in CV, 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 in the market 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 in the market, 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 the market and are healthy.

It follows that the probability that a healthy vulnerable agent in the market is *not* infected in a given meeting is $h_t^w + (1 - \gamma)(1 - h_t^w)$ and hence[3]

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

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

[3] 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 market 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. It is worth noticing, however, that the definition of equilibrium must include 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 $x_0^{(1)} = ... = x_0^{(D)} = 0$, 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.

Assume that w_t and e_t are 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.

To derive more implications of the model one must turn to particular 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.[4] 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_0h_0) equal to $1 - 10^{-5}$.

Parameters (γ, ρ^m, ρ^n) determine the rate of transmission of the virus at the market 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,[5] that implies a distribution of contacts between market 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.[6] Consistently with this finding, we set $\kappa = \frac{1}{5.5}$.

[4] The CDC - Centers for Disease Control and Protection - see here
[5] According to the US Census, the average household size is 2.53 people
[6] 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 (D). Similar to existing literature on this matter,[7] we assume D 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 $e_t = e = w * c$ where c is equal to 0.38. Thus, in the baseline scenario, the benefit of staying at home is 38% of the reward of market participation. 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.[8] Parameter q is the probability that in each period, a vulnerable is selected and forced to go to the market. 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 M 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 M as the discounted value of the utility of the resulting constant. In the benchmark calibration,

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

[8] 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.[9] The upper left hand panel shows that all agents spend every day in the market, in spite of the fact that the probability of contagion is higher there than at home. As a consequence, in this case the model behaves just as if we had assumed that agents had no choice between market and 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 the market, 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 in the market, and more so as the infection rate goes up. About two months into the pandemic, the infection rate peaks, and the fraction of time in the market 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

[9]Figures are collected at the end of the paper.

active cases 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 full market participation, 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 full market participation is suboptimal, so that decision makers reduce their time in the market.

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 market 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 in the market 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 in the market, 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 the market gradually. This process is reversed, however, at about day 130, when market participation drops for a couple of weeks. Finally, at about day 150, decision makers return fully to the market.

This evolution reflects the interaction between the dynamics of infection, individual decision, and the financial incentives implicit in the CARES Act. The key aspect of the Act is that the financial reward to market 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 is the market is quite small (see upper middle panel).

The fact that decision makers increase slightly their participation at markets 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 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 calibrated model, all infection occurs in the market 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 go to the market) are exposed at the market. Hence the first wave is relatively small.

As that wave subsides, people start returning to the market, 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 reduce their time in the market. But then the financial incentives of the CARES Act expire. The consequence is that vulnerables go back to the market fully. 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.

Figure 7 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 in the market, 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 market participation fell during that period, as shown in the upper left panel. When the CARES incentive expire, however, the value of market participation jumps up, prompting vulnerable decision makers to abruptly return to the market. This, in turn, caused an increase in infection rates, and a second, bigger wave of the pandemic.

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 terms of our model, social distancing reduces the number of close contacts at the market (lower ρ^m), while masks reduce the probability of transmission per contact with an asymptomatic (lower γ).

In epidemiological models, these policies would imply a lower and later peak of active cases and infection rate 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. By reducing the probability of infection during market activities, these policies have two opposite effects on behavior. On one hand, a lower probability of infection reduces net infection risk which drives agents to spend more time at markets. On the other hand, an expected path with a lower probability of infection increases the value of being vulnerable relative to being hospitalized, thus, increasing market's expected marginal cost and reducing incentives to go to the market today.

To illustrate the different potential outcomes of these policies in our model, we present in this section three different exercises where we analyze the effect of social distancing.[10] First, we run the benchmark calibration (log utility) with the effectiveness of social distancing, given by the number of close contacts at the market, is set at three levels: 50 percent (the benchmark level), 75 percent, and 95 percent.

As shown by Figure 8, regardless of the effectiveness level, every social distancing scenario drives people to allocate their whole day to market activities, just as in the benchmark case with log utility. Consequently, by reducing the pace of transmission, these policies have a positive effect on the dynamics of the virus with lower peaks and fewer deaths.

[10] The implications of mandatory masks are qualitatively not different from social distancing. Thus, for brevity, we present the results for social distancing only.

Having said this, Figure 9 shows the opposite effects that social distancing has on incentives. Panels (D) and (E) of this Figure show that greater effectiveness in social distancing implies a lower net infection risk peak and, at the same time, a higher expected marginal cost during the initial phase of the pandemic. With the benchmark calibration, the effect on infection risk is greater than the effect on the expected marginal cost, and, thus, the result of going to the market.

In a second exercise, we modify the benchmark calibration by making M (the utility cost of death) five times larger and, again, we implement social distancing policies with three different effectiveness levels. As seen before, a greater M increases the incentive to stay home because it reduces the *value* of being hospitalized.

Under this new calibration, Figure 10 shows that, while in the "m=5" scenario people allocate their time to market activities, with social distancing at 75% and 95%, people maximize by staying some time at home. These choices, albeit at first counter intuitive, result from policies that fail to reduce sufficiently current infection risk to offset the increase in the expected marginal cost (Figure 11).[11]

This second exercise underscores that the decision of going to markets or staying at home depends on this double relative, discussed deeper by CV, and not on the absolute values of probabilities of infection. Moreover, in this case, maximizing behavior complements the goal of social distancing (Figure 10). That is, not only does a lower number of close contacts reduce the severity of the epidemic, but by staying at home the positive effects are greater.[12]

However, under different environments, social distancing and maximizing behavior can actually offset each other. Our third exercise illustrates this case. We ran the same three policies as before. The difference is that we do it in an initial environment where agents optimally reduce their time at the market with no social distancing policy.[13]

[11] This is not the case of a social distancing policy, effective at 50%, because the net infection risk becomes close to zero

[12] For example, with a 95% effectiveness, accumulated deaths in the first exercise climbed up to 1% of the initial population while in the second exercise they were 0.7% of initial population

[13] We use as the new benchmark the scenario with $m = 10$

Moreover, we assume that the introduction of the social distancing policy is unexpected by agents, and it is done at day 40 after the arrival of the virus to the community. Figure 12 shows the results of this exercise for time at market, active cases and accumulated deaths. Three comments are worth making at this point.

First, all social distancing policies drive agents to increase their time at the market, relative to the "m=10" scenario, at least in the short run. Moreover, the greater the effectiveness of the policy, the greater is the increase of market activities.

Second, 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 what is observed for scenario with "m=10". Thus, at least in this calibration, agents don't increase their market activities enough to get a transmission of the virus above the pre-policy state.

Third, a basic SIR model ignores that, after the implementation of a social distancing like policy, agents reevaluate their behavior, and, by doing so, these models overestimate the effect of such policy on epidemiological results. Figure 13 shows that, for every level of efficiency, the peak of active cases would be lower when the model ignores that agents optimize after the introduction of the policy (dotted lines). In our model, by spending more time at market activities, the effect of social distancing is partially offset, and, thus, the number of active cases falls less (dash lines). In the case of a 95% social distancing efficiency, the effect, in equilibrium, of the policy is minimal.

7 A Social Planning Problem

7.1 Externalities and Socially Optimal Outcomes

Our model yields the sequence $\{p_t\}$ of individual decisions an an equilibrium outcome. In equilibrium, vulnerable decision makers choose p_t balancing relative costs and benefits, which are partly given by the development of the virus and contagion probabilities, which individual decision makers take as given. On the other hand, contagion probabilities depend on the distribution of people between market and home, 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 (and, therefore, of all the agents in this economy). The problem can be written in recursive form as follows: if U_t denotes the period social utility, it must be given by:

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

where, for convenience, we have defined $J_t = \mu x_{t-1}^{(D)}$ as the number of deaths in period t. The first term in the sum in the right hand side is the utility of vulnerables, which depends on p_t. The other terms gather the utility of recovered and hospitalized people, minus the cost of deaths. Note that we denote welfare cost of deaths as M_s, instead of M, to signal that this can potentially be different than the individual cost of death.

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 the period t choice variables are p_t and the probability of infection of a healthy vulnerable, which is 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 sheds 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:

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

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 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 population, 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 14.[14]The planning solution differs from the equilibrium outcome (which, remember, is also the outcome of the SIR model) in substantial ways. For the first forty days, the planner allows for full market participation, 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 to market. Full return to market is not observed until after 260 days since the onset of the pandemic.

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

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

optimal to shutdown market activity too early.

Likewise, the planning solution underscores that, after an initial strict lockdown phase, a rapid return to market 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 *controls* the economy's infection rate.

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 planner solution reduces market 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.

7.2 The Sustainability of the Planning Solution

The planner's problem can be used to analyze the sustainability and compliance of lockdowns. Levy-Yeyati et al. (2020) find for 120 countries that lockdown compliance decreases over time, with a stronger trend in economies with higher levels of labor precariousness. This occurs in economies where most jobs cannot be done from home nor there is a strong safety net. Thus, in terms of our model, lockdown compliance should be lower with a greater wedge between w and e.

To study compliance, we assume that planner's solution is the lockdown policy that the government of this economy wants to implement. Then, we ask what is the optimal individual behavior for an agent that believes that infection probabilities are given by the lockdown policy.

Figure 15 presents three series[15] of time allocated to market activities for two different calibrations.[16] Panel (A), which is supposed to represent an advanced economy, has the same sequence of home reward as the benchmark calibration. Meanwhile, the developing economy scenario consists of a home reward sequence is that is 5% of the benchmark calibration for the whole simulation. This difference is supposed to capture the greater labor precariousness in

[15] Decentralized equilibrium, optimal planner, and best response to optimal planner

[16] Both calibrations have the utility cost of dying five times larger than in the benchmark calibration

developing countries.

In the advanced economy, compliance is high since, during the first 60 days, individuals best response follows the planner's solution. However, after that day, an individual maximizes by returning to market immediately while the lockdown policy returns with an tighten and softening pattern and fully reopens the economy two years after the onset of the pandemic.

Interestingly, optimal individual behavior is to reduce market activity together with the lockdown at the beginning while, in the benchmark case, an individual never reduces its allotted market time. This discrepancy is a consequence that the lockdown reduces the expected path of infection rates which increases in the short run the gap between being vulnerable versus generating symptoms (hospitalized). Thus, at the beginning of the lockdown, the infection risk is relatively too high that is not optimal to run the risk of missing out on future lower infection rates by going to the market. Once current net infection risk falls *enough* due to the lockdown, then, it is optimal to return to market activities.

In the case of the developing economy, incentives to comply with the lockdown are lower given that gap between market and home reward is bigger. This can be seen in Figure 15 where the best response reduces market time about 40% while the planner's solution does so by 90%. The wedge between market and home reward is too big that the greater expected marginal cost of infection risk doesn't compensate sufficiently to drive agents to stay at home for more time. As in Levy-Yeyati et al. (2020), countries with lower ability to transfer jobs to households or a weaker safety net find it harder to comply with lockdown policies.

This discussion provides insight on why in some countries lockdown compliance is greater than in others. We show that our model is capable of replicating Levy-Yeyati et al. (2020) analysis. These results underscore that policy compliance is, at least partly, determined by the double relative of economic incentives. Hence, compliance can be improved both by changing economic incentives.

Finally, a comment about the different paths of the planner's solution between advanced and developing economies. Since in a advanced economy the gap between market and home reward is lower, it is less costly economically to turn on and turn off lockdowns in order to control any possible new waves. This is not the case for a developing country. Reinstating a lockdown is difficult because the economic loss is too high, thus, a social planner is willing to incur in more deaths (0.02 percentage points more).

8 Relation to the Literature and Final Remarks

To be Written

9 References

- 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.

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

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

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

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

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

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

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

- 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.

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

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

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

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

- 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.

- 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

- 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.

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

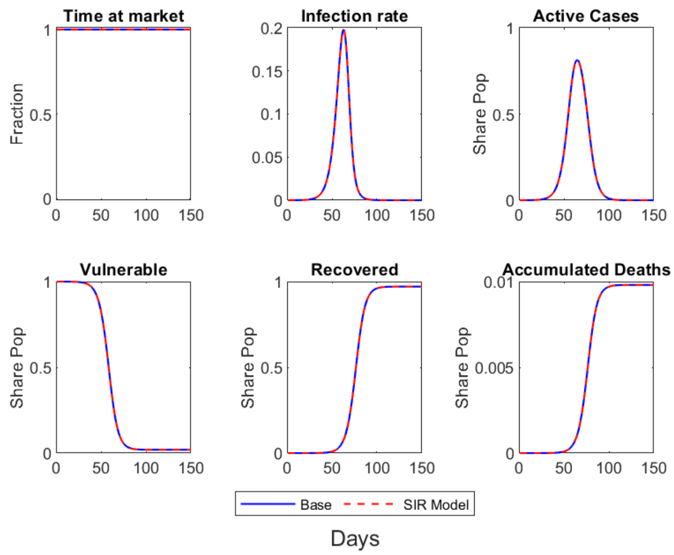

Figure 1: Model (log utility) versus SIR

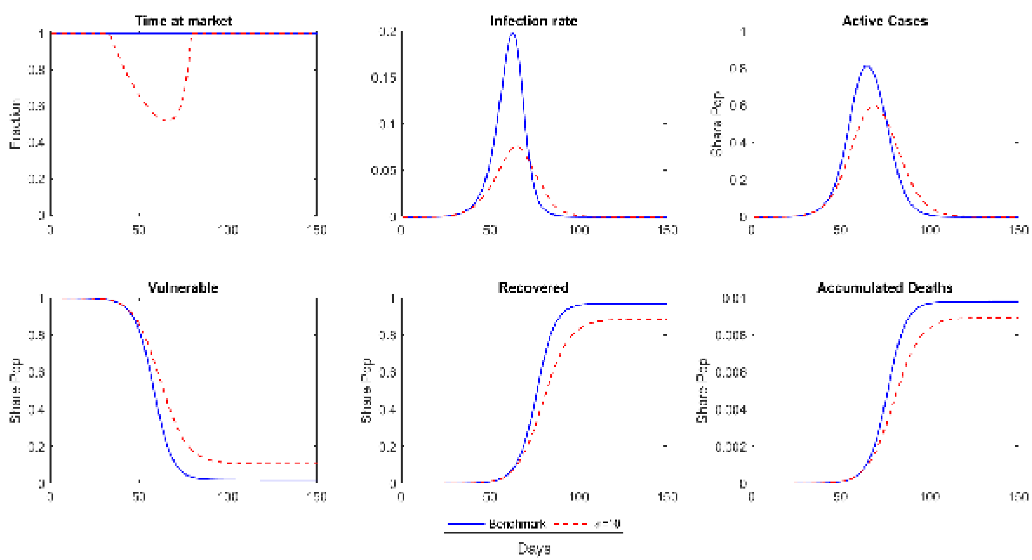

Figure 2: Model ($\sigma = 10$) vs SIR

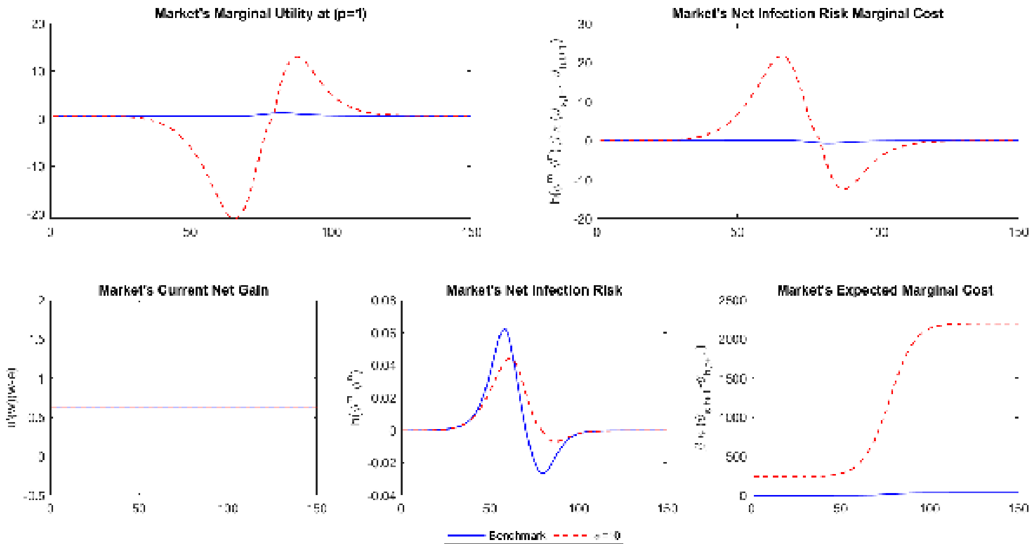

Figure 3: Incentives for Decision Making ($\sigma = 10$)

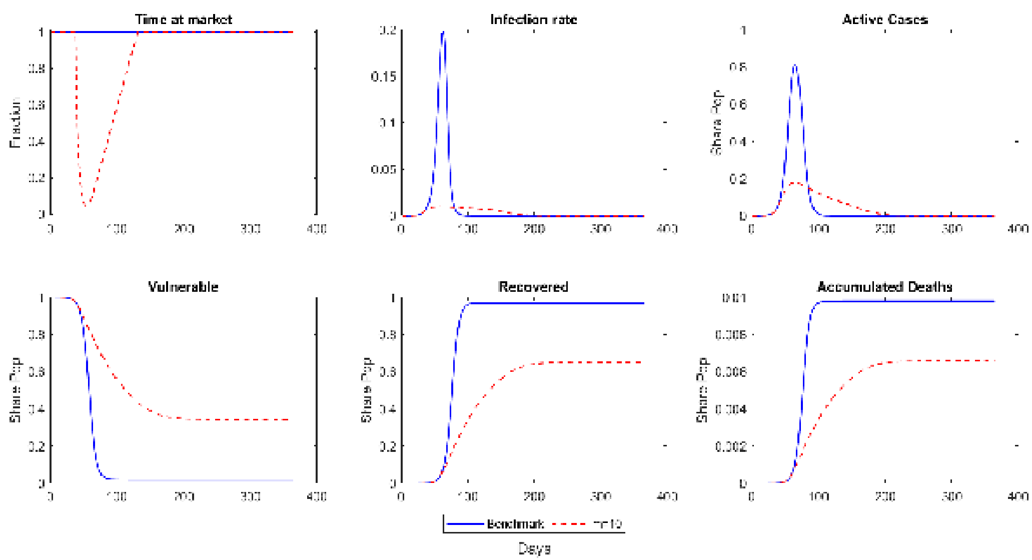

Figure 4: Bigger Fear of Death

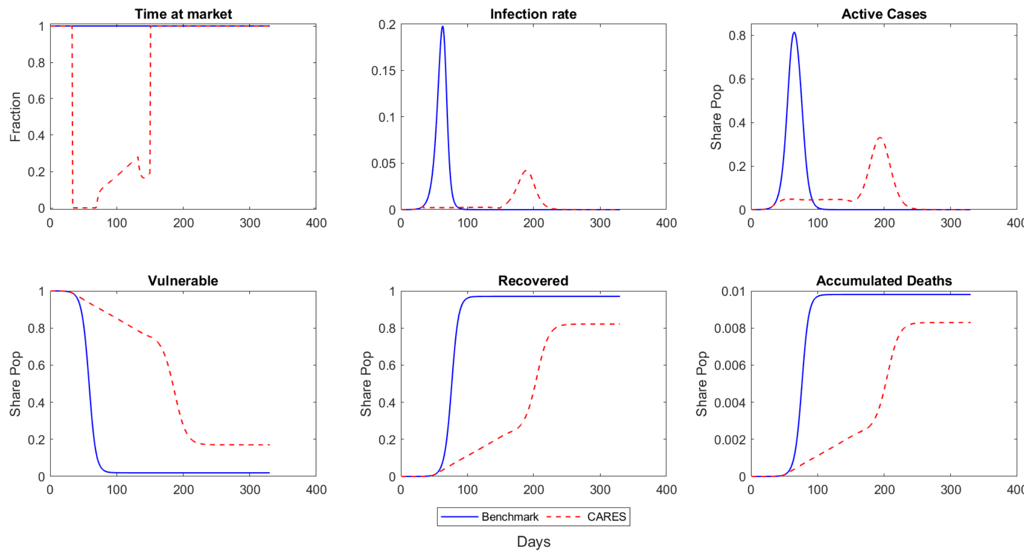

Figure 5: CARES and Multiple Waves

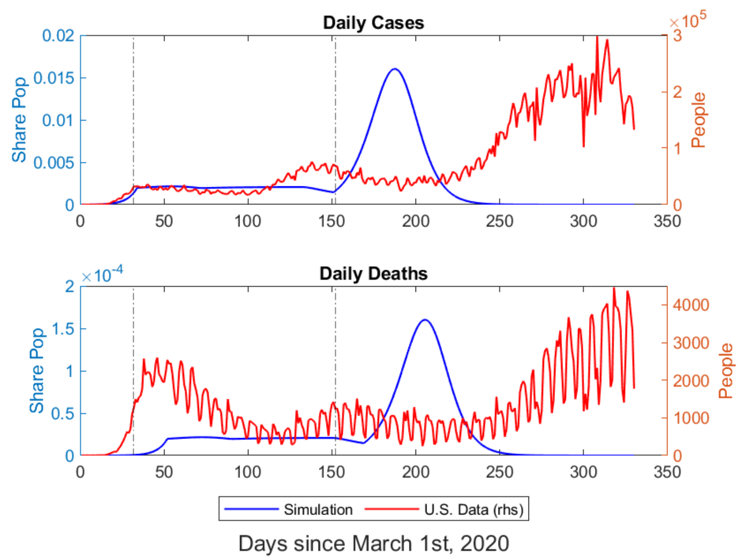

Figure 6: Cases and Deaths: Model and US data

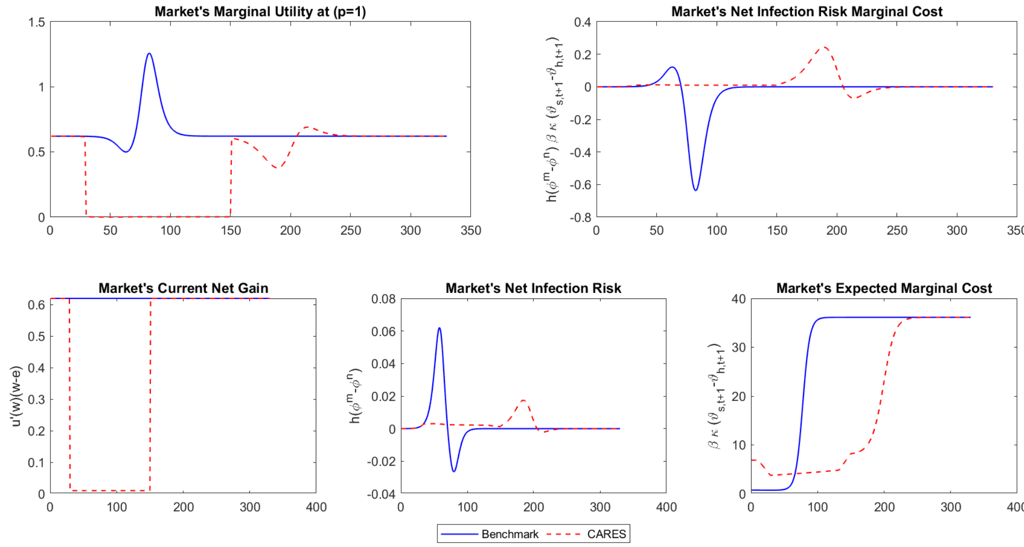

Figure 7: Incentives in the CARES Act

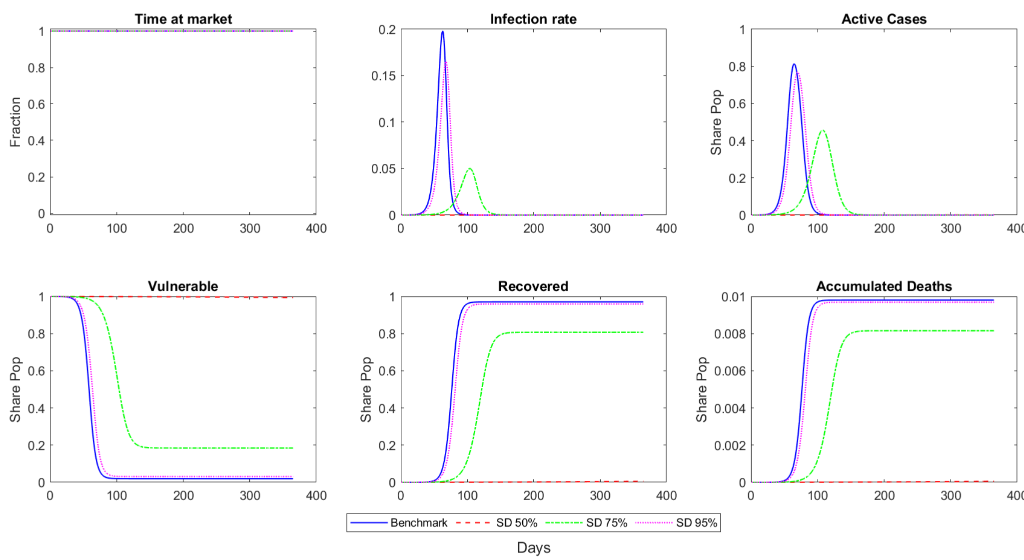

Figure 8: Social Distancing

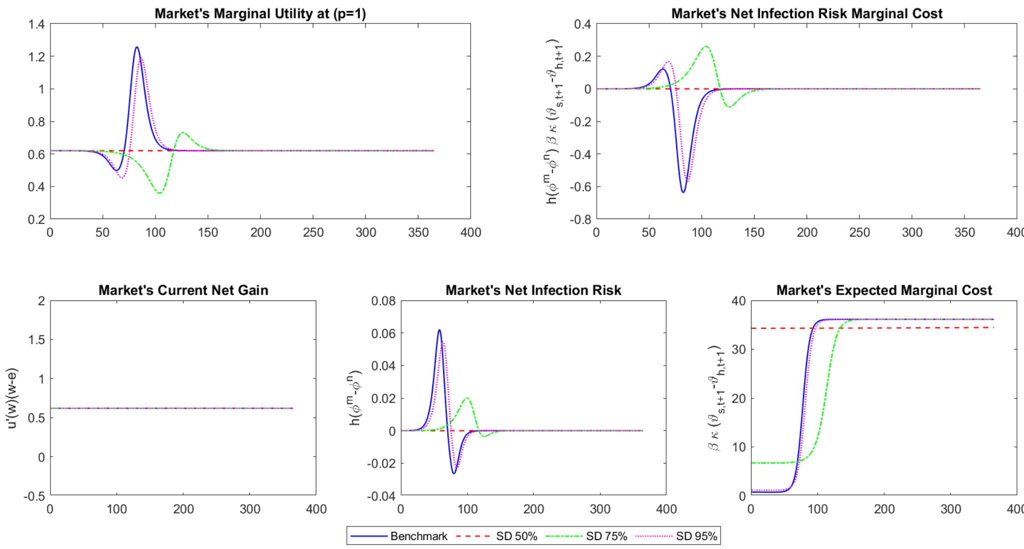

Figure 9: Social Distancing: Incentives

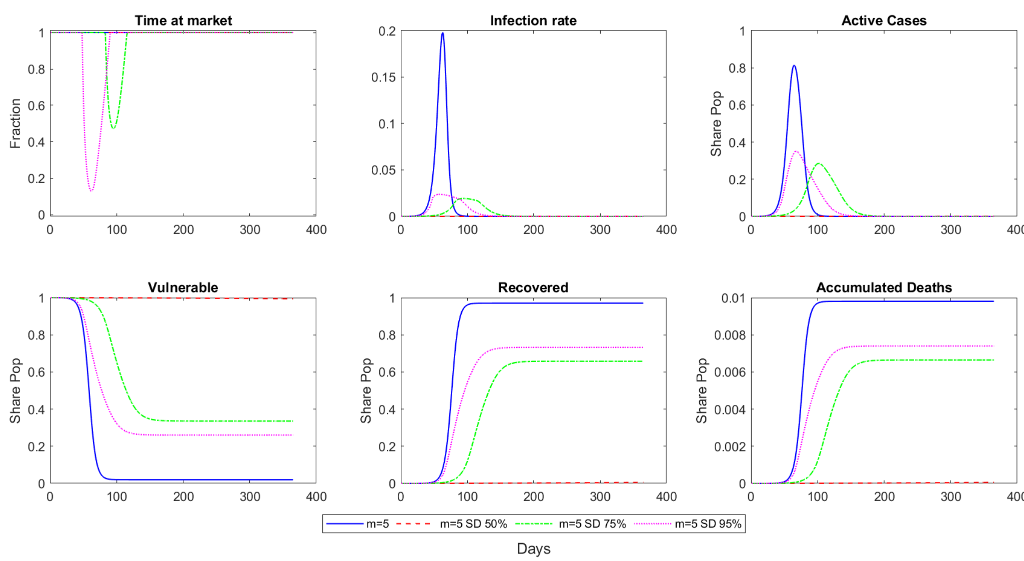

Figure 10: Social Distancing ($m = 5$)

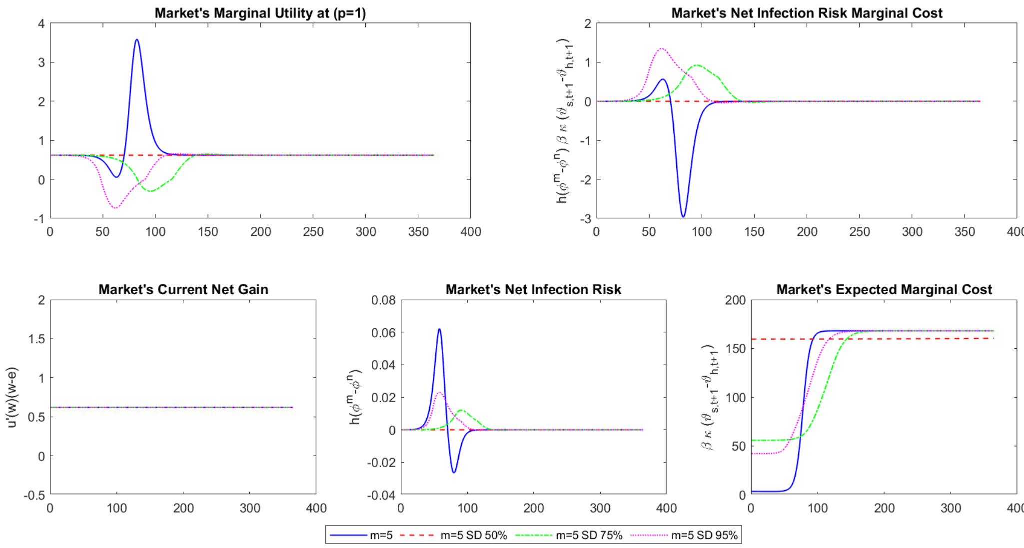

Figure 11: Social Distancing and Incentives ($m = 5$)

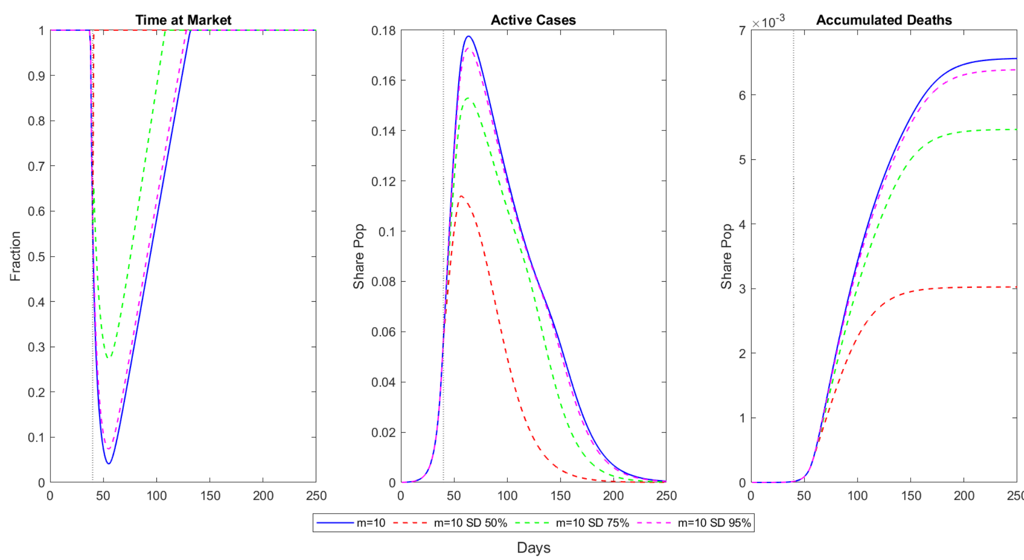

Figure 12: Social Distancing ($m = 10$)

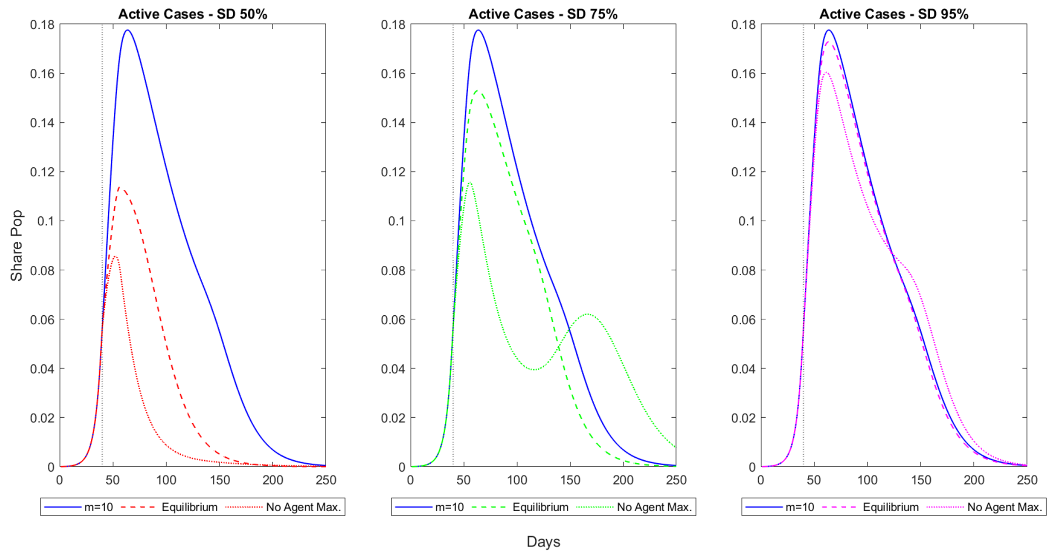

Figure 13: Social Distancing vs SIR

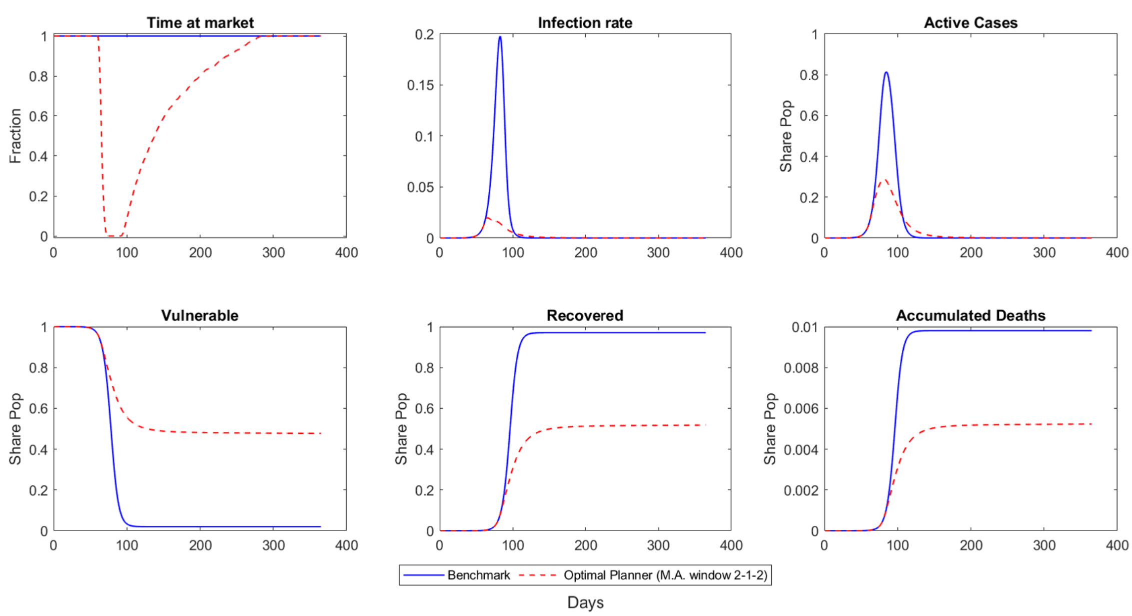

Figure 14: Equilibrium vs Socially Optimal Plan

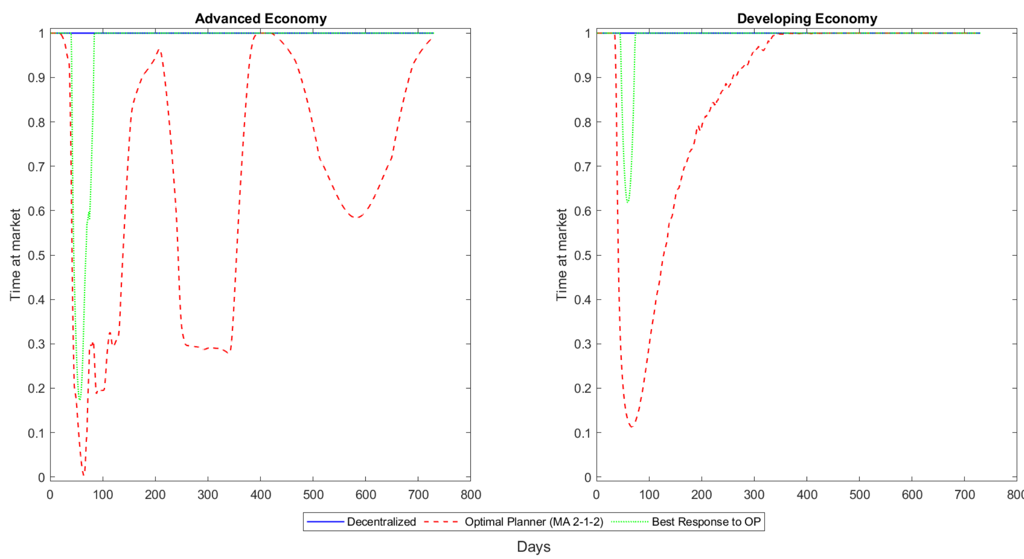

Figure 15: Lockdown and Compliance