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

Roberto Chang[†] Humberto Martínez[‡] Andrés Velasco[§]

First Version: March 2022

This Version: September 2025

Abstract

In a pandemic, the evolution of contagion depends on choices that individuals make in response to incentives, including economic incentives. We develop this idea in a dynamic model of pandemic contagion embedded in an economy where individuals choose whether to work in an outside location or stay at home. Infection rates depend on how many people decide to go to work. In contrast with conventional SIR models, in our model individuals choose locations taking into account the pay gap between outside work and staying at home, as well as the risk of contagion and associated future economic and health-related losses. In equilibrium, the dynamics of contagion are influenced by factors that have no relevance in SIR models, including utility functions, expectations, and, more importantly, economic policies. We explore several novel implications of the links between pandemics and incentives for contagion dynamics and policy.

*We thank Drew Harris, Juan M. Sánchez, and seminar participants at Rutgers University, Banco Central de Chile, the Atlanta Fed, the 4th International Macro-Finance and Sovereign Debt Workshop, and the 2021 LACEA-LAMES conference 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. Humberto Martínez gratefully acknowledges research support to the Chilean Comisión Nacional de Investigación Científica y Tecnológica through project Fondecyt de Iniciación en Investigación 2024 No. 11241286.

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

[‡]Universidad de Chile - FEN hmartinezb@fen.uchile.cl

[§]London School of Economics and Political Science, and CEPR. A.Velasco1@lse.ac.uk

1 Introduction

The Covid-19 pandemic was the biggest economic and social shock in recent history. At the onset of the pandemic, the public policy strategy of most governments worldwide was clear: send people home to slow the transmission of the virus, in the hope of preventing the collapse of health care systems and gaining enough time to expand the capacity of health care services. *"Flattening the curve"* became the guiding motto.[1] Economic policies were deployed to mitigate the side effects of public health policies, rather than to address the evolution of the pandemic itself.

This policy approach was consistent with dominant models of pandemics, especially the Susceptible-Infected-Recovered (SIR) model, in which the evolution of the pandemic is determined by fixed transmission and recovery rates only. In such models economic incentives play no role and policy makers are stuck with choosing between lives and livelihoods with no hope of affecting the terms of the tradeoff.

But while the initial appearance of a virus may be exogenous, the speed of virus propagation is not. The dynamics of a pandemic are the result of human decisions and, hence, can be influenced by economic policies. In particular, the consequences of health policy mandates, such as stay-at-home guidelines, lockdowns, and social distancing, depend on the decisions of individuals who weigh the costs and benefits of compliance. In turn, individual choices respond to both health and economic incentives, including foregone wages, the extent of government financial support for those who stay at home, and expectations of the speed of the post-lockdown economic recovery.

Once we recognize that during a pandemic people respond to economic incentives, several questions arise: can we identify policies, other than health directives, that help control a pandemic? How should we evaluate such policies? How do the policies interact?

To start answering these questions, we develop a dynamic model of the interaction between individual decisions and the dynamics of virus transmission, emphasizing the model's implications

[1] The International Monetary Fund titled its April 2020 World Economic Outlook "The Great Lockdown" reflecting the fact that most governments implemented some type of stay at home mandate in their economies

for policy. Crucially, our model is built on first principles, so that, unlike SIR models, policy analysis is immune to the Lucas critique. (Lucas 1976).

Our framework delivers several new lessons that contrast to those of standard epidemiological models. We show that, because SIR-type models ignore the impact of non-pharmaceutical interventions on incentives and expectations, they can overestimate the benefits of such interventions on welfare, infection rates, and the number of deaths. More generally, an important and intuitive externality appears in the model: while individuals take infection rates, those infection rates depend on the choices of individuals. One consequence is that lockdown-type policies that might be optimal from a social planner's perspective are generally not individually optimal. Hence lockdowns can be difficult to implement without sufficient enforcement or other targeted incentives.

In our model, expectations and forward-looking behavior play a key role. This is in contrast with SIR models and implies, in particular, that multiple equilibria are possible, in which case, policy can play an important role as a coordinating device. Finally, we show that public health directives are not the only game in town at the government's disposal: well-tailored economic policies, such as temporary fiscal transfers, can improve both health outcomes and economic outcomes in a pandemic. In general, our results underscore that policymakers who are aware of incentives in the epidemiological context can influence pandemic dynamics in ways that policies based on fixed-parameter SIR models cannot capture or anticipate.

In our model, and in contrast with SIR models, expectations and forward-looking behavior play a key role. This implies that multiple equilibria are possible, in which case policy can play an important role as a coordinating device. Finally, we show that public health directives are not the only tool at the government's disposal: well-tailored economic policies, such as temporary fiscal transfers, can improve both health outcomes and economic outcomes in the course of a pandemic. Our results underscore that policymakers that are aware of incentives in the epidemiological context can influence pandemic dynamics in ways that policies based on fixed- parameter SIR models cannot capture or even anticipate.

Our model is built on ideas from Chang and Velasco (2020) extended to an infinite time horizon. Aside from the obvious gain in realism, the extension makes the model directly comparable to the conventional SIR model, as well as more useful for policy analysis. The setting features a population that normally works outside the home. This normal state of affairs is interrupted by the appearance of a contagious virus. During the ensuing pandemic, infected individuals can become sick, in whose case they must stay in a "hospital", where lucky ones recover and unlucky ones die. Other people, referred to as vulnerable, can be infected but not show symptoms of the virus. Vulnerable individuals choose whether to go to work or stay at home. Their decision must take into account that the probability of infection is greater at work than at home, but staying at home is potentially costly in terms of lost income. Hence vulnerable agents' choice hinges on a "double relative": the current payoff from outside work relative to home income must be compared to the expected future value of remaining vulnerable relative to the cost of becoming infected.

In evaluating this 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. Equilibrium is defined in the natural way, requiring that work-stay-home decisions and infection probabilities be simultaneously determined.

We are able to derive some interesting results analytically. To derive additional insights we simulate equilibrium paths numerically with parameters calibrated to match U.S. data. Resorting to numerical methods is necessary since the model is highly nonlinear and nonstationary.

For our benchmark calibrations, the model yields a pattern of infection with the hump shape characteristic of a pandemic, similar to what a standard SIR model would yield. But in our model infection dynamics reflect an equilibrium in which people choose to fully work outside, taking into account incentives. One implication is that, in equilibrium, health and economic outcomes depend on fundamentals that would have no bearing in other models. For instance, we show that, *ceteris paribus*, higher risk aversion or a bigger cost of death induce agents to spend more time at home during the pandemic, reducing infection rates and the prevalence of illness and deaths.

Our model is perhaps most useful for the analysis of welfare and policy. In the model, individuals base their choices of location on their own probabilities of contagion, failing to take into account the impact of those choices on the economy-wide dynamics of contagion. Hence, as mentioned already, the model incorporates an externality that can potentially rationalize government intervention.

Assuming that the social objective is the discounted expected welfare of the different epidemiological groups, each weighted by their population size, we contrast the decentralized outcome to the optimal social planning solution. In the benchmark calibration, the planning solution prescribes an initial drastic reduction in the amount of time outside the home, followed by a gradual return to outside work. The solution thus resembles a lockdown and stands in contrast to the decentralized equilibrium, in which people continue to work normally outside the home. Remarkably, the social-planning policy reduces infection but does not eliminate it completely.

We can not only compare dynamics in equilibrium versus an optimal policy, but also identify incentives that can make people comply with the optimal path. Specifically, we can ask what is an individual's best response if the government implements the optimal lockdown. We argue that, if an individual believes that others will comply with the lockdown, the individual's best response is to return to work full time. This reflects the externality mentioned before and suggests that 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, compliance to a lockdown will be, at best, short-lived.[2]

One interesting finding is that *animal spirits* can be crucial in a pandemic: for some parameterizations, multiple equilibria emerge in our model. In those cases, a SIR-type equilibrium, with agents fully working outside the home and high rates of infection and death, coexists with another

[2] These results are consistent with Levy-Yeyati and Sartorio (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

equilibrium in which precautionary behavior results in less time outside the home, lower infection rates, and fewer deaths. In the latter equilibrium individual decision makers cut 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 sufficiently over time so that the future value of not being infected is high. Multiple equilibria thus reflect the importance of forward-looking behavior, in sharp contrast to SIR models, which are only backward-looking.

When multiple equilibria exist, policy can be highly effective acting as a coordinating device. If a SIR-type equilibrium coexists with a precautionary equilibrium, an appropriate mandate to limit time outside the home eliminates the bad equilibrium and is politically costless: in the surviving precautionary equilibrium, decision-makers find it individually optimal to follow the government's directive. Likewise, economic incentive policies can act as coordinating mechanisms, at least in principle.

To emphasize the usefulness of our model for policy analysis, we discuss two more specific applications. One is a *social distancing* policy, modeled as a reduction in the number of contacts between individuals working outside the home. Conditional on people's choices, social distancing reduces current and future infection probabilities, with two opposite effects on individual incentives. Lower infection risk directly increases today's 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 choices is not the absolute value of infection rates but the "double relative" explained earlier.

As a consequence, we find that social distancing policies can increase the share of the population going to work for one set of parameter values, and reduce it for another. Perhaps importantly, we show that, since SIR models ignore that agents reevaluate their behavior after the implementation of a social distancing policy, these models tend to overestimate the impact of such policy on health outcomes.

Our final application involves an economic policy: the implications of a fiscal policy package similar to the U.S. 2020 CARES Act. Specifically, we feed into the model an increase in unemployment benefits that virtually eliminates the difference between market income and home income, but only for a limited time period. In the resulting equilibrium, the pandemic comes in two waves, and the intuition for this should be clear. While they last, fiscal transfers reduce the relative payoff of working outside the home. These incentives induce individuals to reduce time working outside the home, which helps limiting contagion, contributing to a reduction in infections after an initial peak. When the policy expires, the payoff of outside work relative to staying home jumps back to its usual level. Individuals then return to work outside the home, helping start a second wave of the pandemic. In spite of multiple waves, the fiscal package reduces both infection rates and accumulated deaths relative to the no policy scenario.

Our analysis confirms that changes in economic policies, which would not have an impact in SIR models, matter for the decisions of individuals and, more importantly, for the dynamics of contagion and virus transmission. Furthermore, we show that the aspects of the fiscal package that matter for the evolution of the pandemic are those affecting individual incentives and individual decision-making. That is, policies that do not change relative incentives have no impact on dynamics. This is the case, for example, of universal cash transfers to households, which in our model have no impact on the double relative and, hence, on equilibria.

This paper is a contribution to the literature, motivated by the Covid-19 pandemic, that blends models from economics and epidemiology. 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 other policies, including economic policies, 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.[3] There are exceptions, however, to which our paper is related. A number of studies, including Garibaldi et al. (2020), Rachel (2022), and Toxvaerd (2020), 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 mandates. By contrast, we emphasize the impact of economic policy incentives and of policy interventions.

Eichenbaum et al. (2021) and Jones et al. (2021) develop dynamic models that postulate SIR-type equations, which depend on economic activities such as consumption and hours worked. As in our paper, individual agents understand that their decisions about consumption and labor supply have implications for their exposure to contagion, so market incentives such as wages matter for the dynamics of infection. But our paper differs from those of 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

[3]Brodeur et al. (2021) provide a survey of the economic literature around Covid-19.

taxes, because they wash out in the double relative, have no impact on individual choices in our model, and therefore have no effect on contagion dynamics.

Whether these differences are important hinges on the specific objectives of the papers. The main focus of Eichenbaum et al. and Jones et al. is to describe and quantify 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.[4]

The rest of the paper proceeds as follows. Section 2 presents the model, specifying the economic environment, the impact of a virus, the interplay between individual decisions and infection rates, and the definition of equilibrium. Section 3 proposes a numerical calibration of the model and investigates model dynamics. The social planning problem is the focus of Section 4, while multiple equilibria and the role of forward-looking behavior are discussed in section 5. Section 6 discusses social distancing and fiscal policy. Section 7 closes with conclusions and suggestions for further research.

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

2 Model

2.1 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 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.[5] 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:

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

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

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

$\sigma > 0:$

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

2.2 The 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 each 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 $s_0 = 1$.

2.2.1 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

$$\bar{\phi}_t = \bar{\phi}_t(p_t) \equiv p_t\phi_t^m + (1-p_t)\phi_t^n \tag{2}$$

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. Letting $x_{t+1}^{(1)}$ denote the number of vulnerables that are interned in the hospital at the end of period t:

$$x_{t+1}^{(1)} = \kappa\{(1-h_t) + h_t[q\phi_t^m + (1-q)\bar{\phi}_t]\}s_t \tag{3}$$

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

$$s_{t+1} = s_t - x_{t+1}^{(1)} \tag{4}$$

and the number of healthy vulnerables at the start of period $t+1$ is:

$$h_{t+1}s_{t+1} = h_t s_t \left\{1 - [q\phi_t^m + (1-q)\bar{\phi}_t]\right\} \tag{5}$$

2.2.2 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 until they spend time in the hospital, and 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:

$$v_{ht} = \sum_{j=1}^{H} \beta^{j-1} u(e_{t+j-1}) + \beta^{H} \left[(1-\mu) v_{z,t+H} - \mu D\right] \tag{6}$$

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

$$z_{t+1} = z_t + (1-\mu) x_t^{(H)} \tag{7}$$

The law of motion of $x_t^{(1)}$ was given in (3). In turn, by definition,

$$x_{t+1}^{(i)} = x_t^{(i-1)}, \quad i = 2, 3, ...H \tag{8}$$

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, which is crucial.

2.2.3 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:

$$\begin{aligned} v_{qt} &= u(w_t) + \beta[\kappa\{(1-h_t) + h_t\phi_t^m\}v_{ht+1} \\ &\quad +(1-\kappa\{(1-h_t)+h_t\phi_t^m\})v_{st+1}] \end{aligned} \tag{9}$$

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 $\bar{\phi}_t = p_t\phi_t^m + (1-p_t)\phi_t^n$. Note that she takes ϕ_t^m and ϕ_t^n as given.

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

$$
\begin{aligned}
v_{dt} &= Max_{0\le p_t\le 1}\, u(c_t) + \beta\{\kappa[(1-h_t) + h_t\bar{\phi}_t]v_{h,t+1} \\
&\quad +[1-\kappa((1-h_t) + h_t\bar{\phi}_t)]v_{s,t+1}\} \\
\text{subject to (2) and } c_t &= p_t w_t + (1-p_t)e_t
\end{aligned} \tag{10}
$$

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

$$
u'(c_t)(w_t - e_t) - \beta\kappa h_t(\phi_t^m - \phi_t^n)(v_{st+1} - v_{ht+1}) \tag{11}
$$

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 probability 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 the future payoff of remaining vulnerable versus going to hospital. The choice problem has an intratemporal and an intertemporal dimension. Expectations then play a crucial role.

Finally, our assumptions imply that

$$v_{st} = qv_{qt} + (1-q)v_{dt} \tag{12}$$

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.

2.2.4 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 *not* meeting an infected individual in a given match is equal to the proportion of healthy agents outside, given by:

$$h_t^w = \frac{[q + (1-q)p_t]h_t s_t + z_t}{[q + (1-q)p_t]s_t + z_t} \tag{13}$$

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[6]

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

The expression is intuitive. An increase in the proportion of infected people in the market (a fall in h_t^w) 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 such as mandating the use of face masks may affect the probability of transmission γ.

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

Analogous reasoning implies that

$$\phi_t^n = 1 - [h_t + (1-\gamma)(1-h_t)]^{\rho^n} \tag{15}$$

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. A central aspect of the model is that 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.

2.2.5 Equilibrium

Equilibrium can be defined in a natural way. The definition includes the dynamics of infection, in addition to individual decisions. We assume perfect foresight. This follows recent literature, but we comment on this modeling choice in the concluding section.

An **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 (5), (4), (8), (3), and (7).
- The value functions satisfy (9), (10), (12), and (1), (6), given $\{\phi_t, h_t\}$
- p_t attains the max in the RHS of the Bellman equation (10)
- ϕ_t^m and ϕ_t^n are given by (14) and (15)

As we have emphasized, our model is similar to existing SIR models, except (crucially) that p_t is given by individual decisions, so that contagion probabilities are 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 (4) and (3) 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.

3 A Benchmark Case

3.1 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. during the COVID-19 pandemic. According to official numbers, by the end of the first week of March 2020, there were about 338 active cases identified in the country.[7] 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,[8] 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 with the COVID-19

[7] The CDC - Centers for Disease Control and Protection - see here

[8] According to the US Census, the average household size is 2.53 people

virus.[9] Consistently with this finding, we set $\kappa = \frac{1}{5.5}$.

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,[10] 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 between 0.39% -1.33%, with a 95% of confidence, COVID's infection fatality rate in China.

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.[11] Parameter q is the probability that in each period, a vulnerable is selected to work outside. Following Alvarez et al. (2021), 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, 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

[9] See Page 11, Final Report of the mission here

[10] See Acemoglu et al. (2021); Eichenbaum et al. (2021); Alvarez et al. (2021); Verity et al. (2020) for example.

[11] See Five Thirty Eight

constant. Our benchmark calibration assumes that agents have log utility preferences.

3.2 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.[12] For comparison, the outcomes of a standard SIR model are also shown, although in this case the two cases deliver indistinguishable outcomes (in fact, each panel displays a solid blue line corresponding to the SIR model and a dashed red line with equilibrium outcomes of our model; they coincide).

As shown in the upper left hand panel, all agents spend all of their time outside their homes in both our model and the SIR model. In our model, this reflects that individuals decide to work outside despite 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 did not have the option to stay home. The model effectively becomes a standard SIR model.

Figure 1: Model (log utility) versus SIR

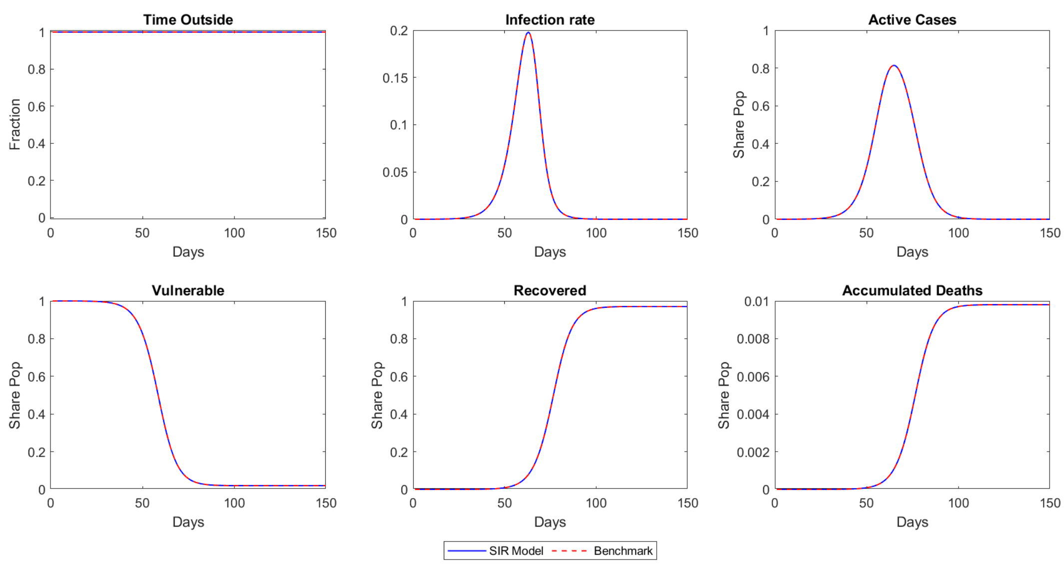

[12]We solve the model implementing a Shooting Algorithm as described in Garibaldi et al. (2020). Code is available upon request.

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, this fact 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 parameter σ equal to ten.

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

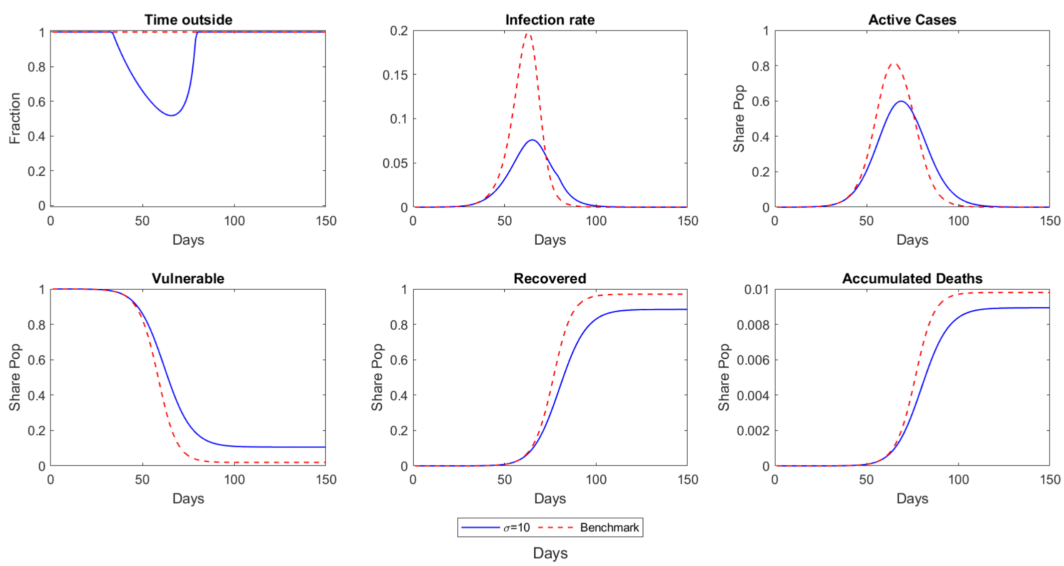

In Figure 2, the case $\sigma = 10$ is displayed by the blue solid line. As in the SIR model, the rate of infection starts accelerating about thirty days after the initial seed. But at that time, and in contrast with the SIR model and the case with log utility, the upper left panel of Figure 2 shows decision making vulnerables reducing their time outside, and more so as the infection rate goes up. About

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

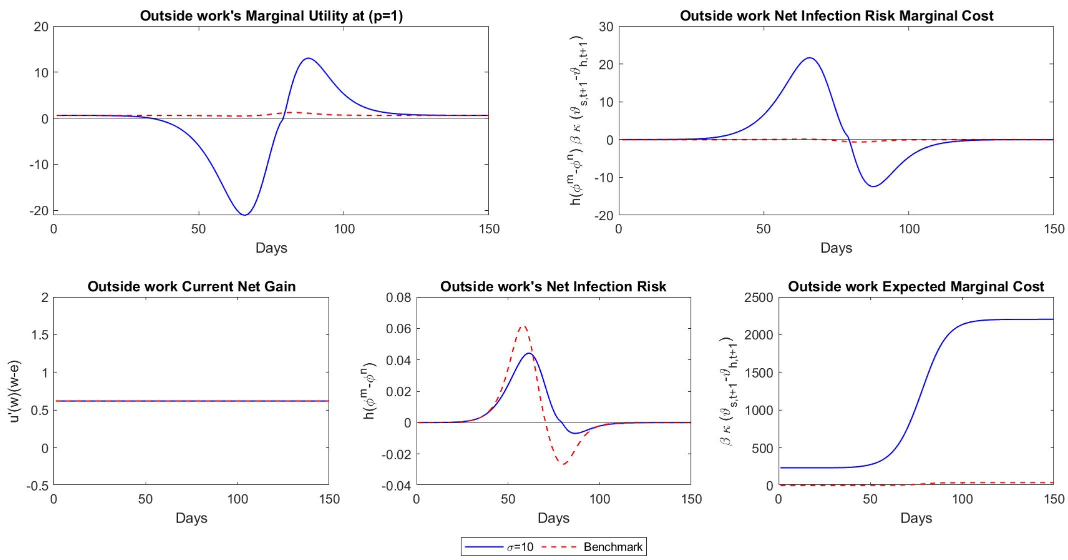

two months into the pandemic, the infection rate peaks, and the fraction of time outside bottoms at around one half. These two variables interact: since people stay home, the infection rate peaks at a lower level than in the SIR model. As a consequence, the number of active cases falls, and the total 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, (11), 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 11. 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.

Figure 4: Model with Bigger Fear of Death

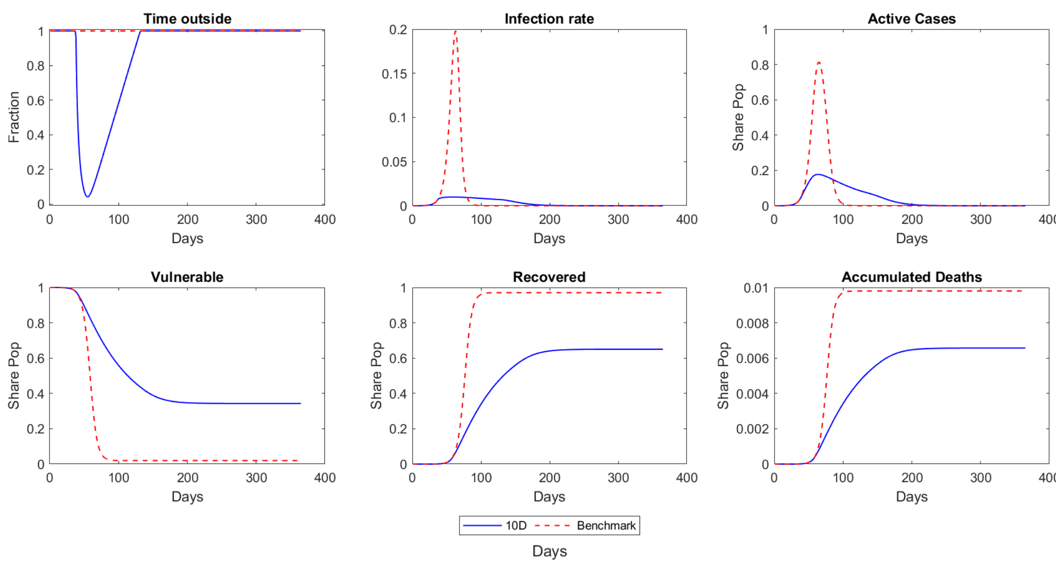

To underscore the crucial role of individual responses to incentives, Figure 4 compares the original calibration, including log utility, against a case in which the utility cost of death is ten times larger. In our model, this means that vulnerable decision makers are more prone to stay home, rather 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 cumulative number of deaths in the long run falls to six tenths of one percent of the population.

4 A Social Planning Problem

In any equilibrium of our model, vulnerable decision makers choose location balancing relative costs and benefits, including those related to the virus and contagion, which individuals take as given. On the other hand, contagion probabilities depend on the distribution of people at home and outside, which is determined by people's locations. 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

$$\Phi_t = (q + (1-q)p_t)\phi_t^m + (1-q)(1-p_t)\phi_t^n \tag{16}$$

together with (13), (14), and (15). Meanwhile, state variables evolve according to rewritten (3), (4), (5), (7), and (8) using (16).

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 (16).

Comparing the preceding expressions against the corresponding expression for individual decision makers (equation (11)), 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

Figure 5: Equilibrium vs Socially Optimal Plan

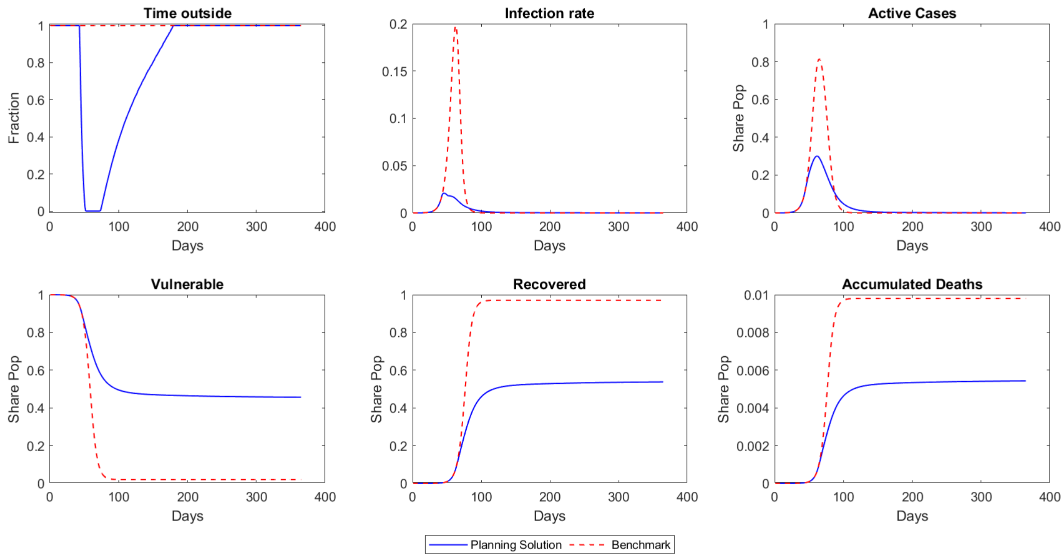

case are illustrated in Figure 5.The planning solution differs from the equilibrium outcome (which, remember, is also the outcome of the SIR model) in substantial ways. For the first forty three days, the planner allows agents to work fully outside, as in the equilibrium model. But then the planner sets p_t to almost zero for about three 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 180 days since the onset of the pandemic.

Notably, the planning solution reduces outside activities only after 43 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 percent 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.

Importantly, the externality we have identified implies that the planning outcome will generally 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.

5 Pandemics and Animal Spirits: Multiplicity of Equilibria

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. To illustrate, 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 3.2, it replicates the outcome of a standard 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, (11) 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

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

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 $\hat{h}_\tau = \tilde{h}_\tau$, $\hat{\phi}^n_\tau = \hat{\phi}^n_\tau$, while $\hat{\phi}^m_\tau \geq \tilde{\phi}^m_t$, since ϕ^m_t 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 $\hat{\phi}^m_\tau \geq \tilde{\phi}^m_t$). 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

Figure 6: Multiple Equilibria: Pandemic Dynamics

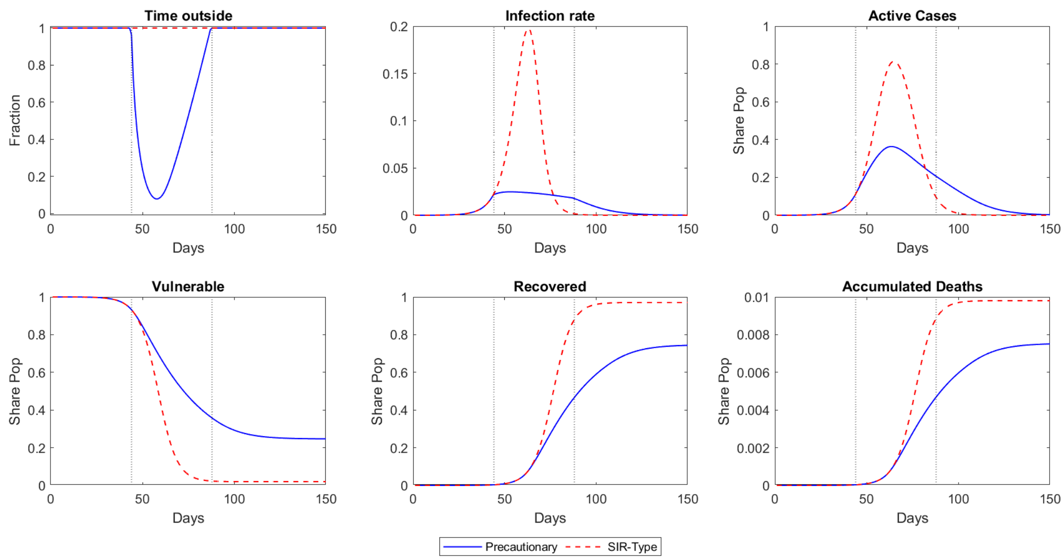

model they must be dynamic ones. Expectations can therefore play a key role.

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 Section 3 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 disappears 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 6 shows the simulation results for both equilibria. In this case, τ is equal to 44; also,

Figure 7: Multiple Equilibria: Individual Incentives

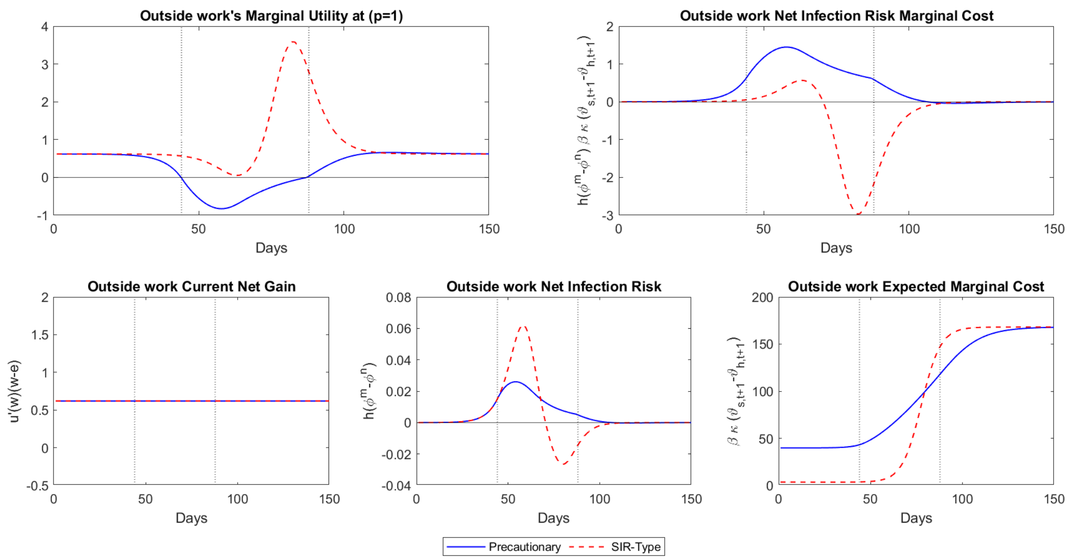

market participation is full in the precautionary equilibrium after 88 days. Thus, the precautionary equilibrium deviates from the SIR-type equilibrium for over a month. As a result, the precautionary equilibrium reaches a steady state with about one third fewer accumulated deaths than in the SIR-type 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 6) 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 remain 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 7 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 novel and 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.

6 Implications for Policy Analysis

We have seen that the incorporation of decision making individuals has important consequences for the analysis of a pandemic and its dynamics. Naturally, the model has several consequences for public policy. We illustrate some of these implications in two cases.

6.1 Non-Pharmaceutical Interventions: Social Distancing

In the recent pandemic, social distancing and mandatory masking were commonly imposed in the hope of limiting the transmission of the virus between people. In our model, social distancing can 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, the two policies reduce the pace of transmission, implying a lower and later peak of active cases. 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 less certain, due to their impact on individual incentives. By reducing the probability of infection associated with outside activities, social distancing and mask wearing have at least two 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. These two effects work in opposite directions.

To illustrate, we analyze the effect of social distancing policies in several alternative cases. We restrict our discussion to social distancing since we have found that the implications of mandatory masks are qualitatively not different. We start by solving the benchmark calibration (log utility) assuming that the effectiveness of social distancing, given by the number of close contacts outside the home, is either high (a 50 percent reduction of ρ^m) or low (a 25 percent reduction).

The main outcomes are displayed in Figure 8. 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. In fact, in the high social distancing scenario, a pandemic never takes off.

Figure 8: Social Distancing

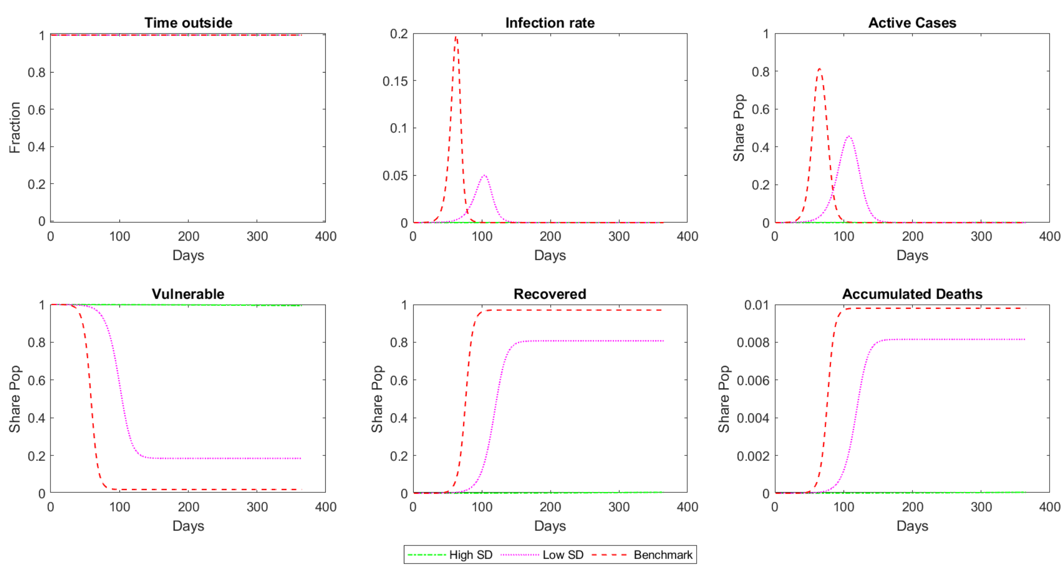

While Figure 8 shows that social distancing does not induce people to reduce their time outside, it does change the incentives that vulnerable decision makers face. This is illustrated in Figure 9. The middle and right panels in the bottom row of the figure show that greater effectiveness in social distancing 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.

Figure 9: Incentives and Social Distancing

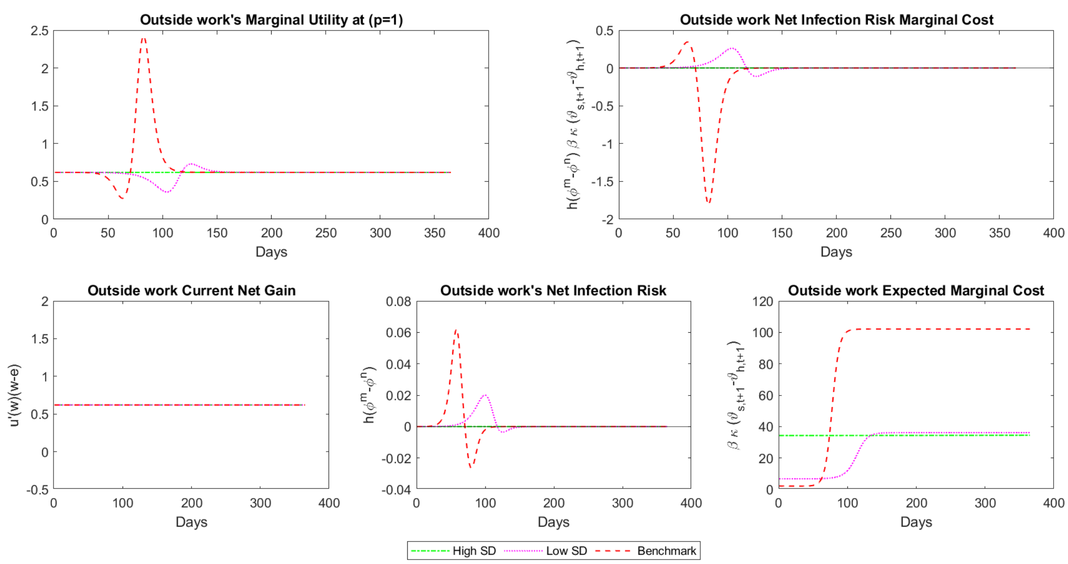

Intuitively, therefore, social distancing might induce decision makers to reduce time outside if the expected marginal cost of infection were greater. To check this conjecture, we modify the previous exercise by making the utility cost of death three times larger than in the benchmark.

As discussed before, a greater D increases the incentive to stay home because it reduces the *value* of being hospitalized. Figure 10 shows our findings. In this "$3D$" scenario, people choose to stay more time at home only if social distancing effectiveness is at 25% (low SD). The result is intuitive. When social distancing is only weakly effective (or not implemented), the incentives effect is not strong enough to induce people to reduce outside work time. Meanwhile, If social distancing is very highly effective, infection risk is driven to negligible levels, so people have less fear to work

Figure 10: Social Distancing With Cost of Death $= 3D$

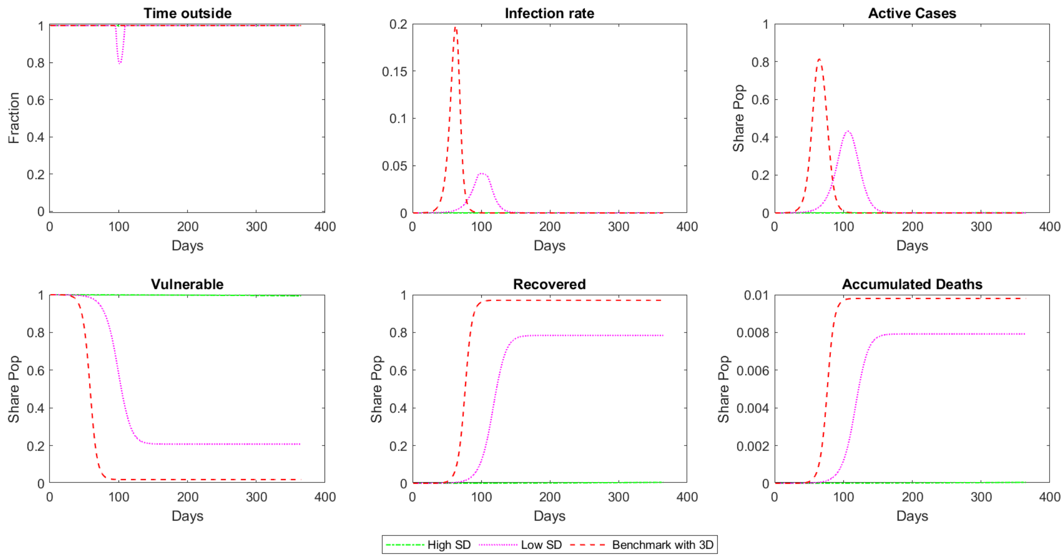

outside.

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 decision making 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 low SD, 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 decision making 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 11 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.

Figure 11: Social Distancing (cost of death= $10D$)

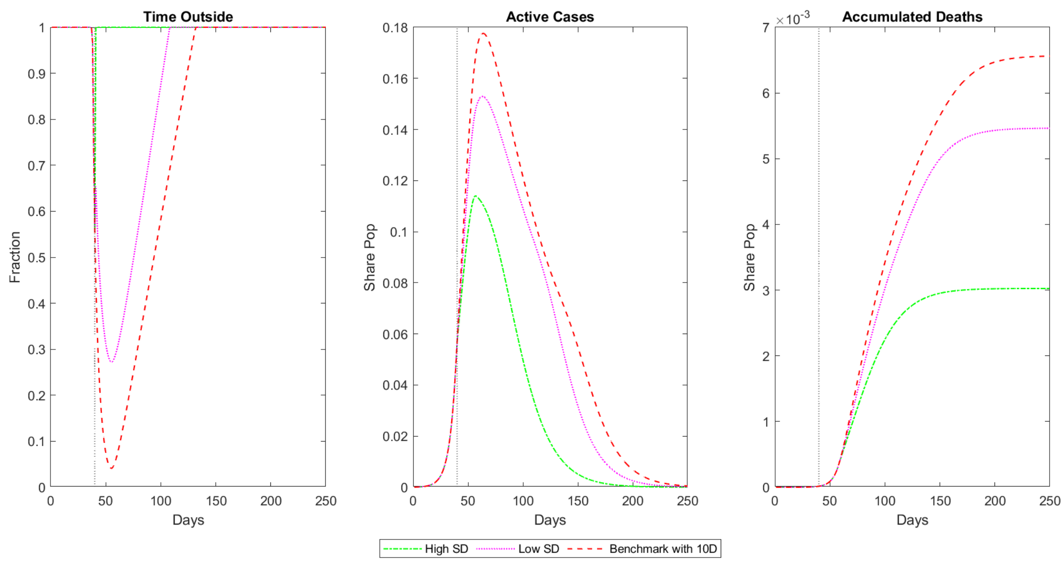

Notably, despite driving agents to spend more time at the market, social distancing policies still manage to produce less severe epidemiological results. This is because of the direct effect of social distancing on infection probabilities. Active cases peaks and accumulated deaths for the two efficiency 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 basic SIR models usually ignore that agents reevaluate their behavior after the implementation of a social distancing policy. By doing so, SIR-type models would overestimate the effect of such policy on epidemiological results. To see this, Figure 12 presents the time outside, evolution of active cases and total deaths when the cost of death is ten times the benchmark, for three different scenarios: our model without social distancing (red dashed line), our model with social distancing (dash-dot green line) and the SIR model type outcome where there is social distancing but agents are not maximizing (blue solid line). The peak of active cases and total accumulated deaths under SIR-type assumptions would be lower. This is, of course, because in our model the effect of social distancing is partly offset by a change in individual behavior (first panel, Figure 12): people respond to the policy by spending more time outside, which limits the impact of the policy on the dynamics of the pandemic.

Figure 12: Social Distancing: vs SIR

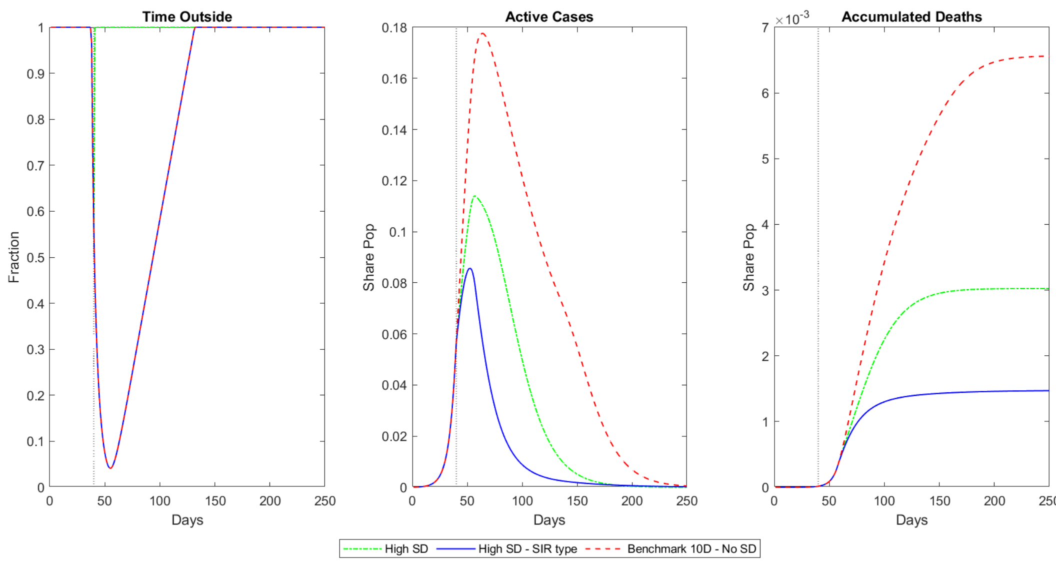

6.2 Fiscal Policy

In March of 2020, the U.S. Congress reached an agreement to provide a economic relief package worth approximately $2 trillion dollars denominated The Coronavirus Aid, Relief, and Economic Security Act (or CARES Act). Within a wide range of different measures, the CARES Act redesigned the unemployment benefits program by including an additional $600 dollars per week to beneficiaries, an increase of 13 weeks in the time window to receive unemployment benefits and, lastly, an expansion on the eligibility criteria to self-employed and gig workers. Importantly, these changes were temporary as the policy was active between April and June of 2020.

To infer the effects of a fiscal policy similar to the CARES Act in our model, we simulate an economy where the calibration of all parameters remain as in the benchmark case except for the schedule of the benefits of staying home. More precisely, we increase e_t from 0.38 to 0.99 between day 30 and day 120, after which e_t returns to its initial value of 0.38. This parametrization in e_t captures both the increase in the unemployment benefits set in the Cares Act as well as the temporary aspect of this policy.[13]

Figure 13: Transitory Fiscal Policy

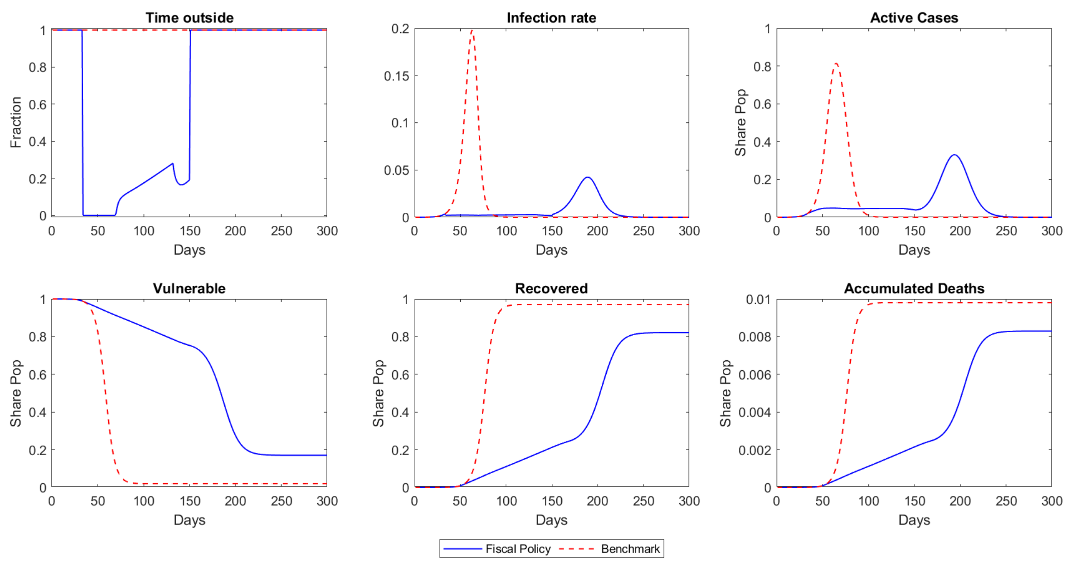

Figure 13 illustrates the implications of this fiscal policy for our model. The upper left panel of this figure shows that, upon the increase in e_t, vulnerable decision makers choose to stay home for an initial period of about a month, after which they start returning to working outside gradually. This 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. The behavioral

[13]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. 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. 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.

response to the fiscal policy contrasts with the benchmark case where individuals do not reduce their time outside. Naturally, the dynamic of the pandemic also differs relative the the benchmark case. 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.

The first wave of infection and active cases is relatively is small since vulnerable decision makers stay at home initially where the number of close contacts is smaller, and only essential workers (vulnerables that must work outside) are exposed. Additionally, this first wave is also flat and remains flat up until day 150 which is beyond the period that the fiscal policy is active.

As this first 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 fiscal policy 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.

This evolution of the pandemic reflects the interaction between the dynamics of infection, individual decisions, and the financial incentives implicit in this fiscal policy. The key aspect of the policy is that the financial reward to outside work relative to staying at home becomes effectively zero for a while. Then, as soon as the policy is implemented, vulnerable decision makers choose to stay at home. Note that this transition is quite abrupt. 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 fiscal incentives 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.

Figure 14 displays how the incentives embedded in the fiscal policy affect the trajectory of the pandemic. The lower left hand panel displays the evolution of the financial incentive to work

Figure 14: Incentives and Transitory Fiscal Policy

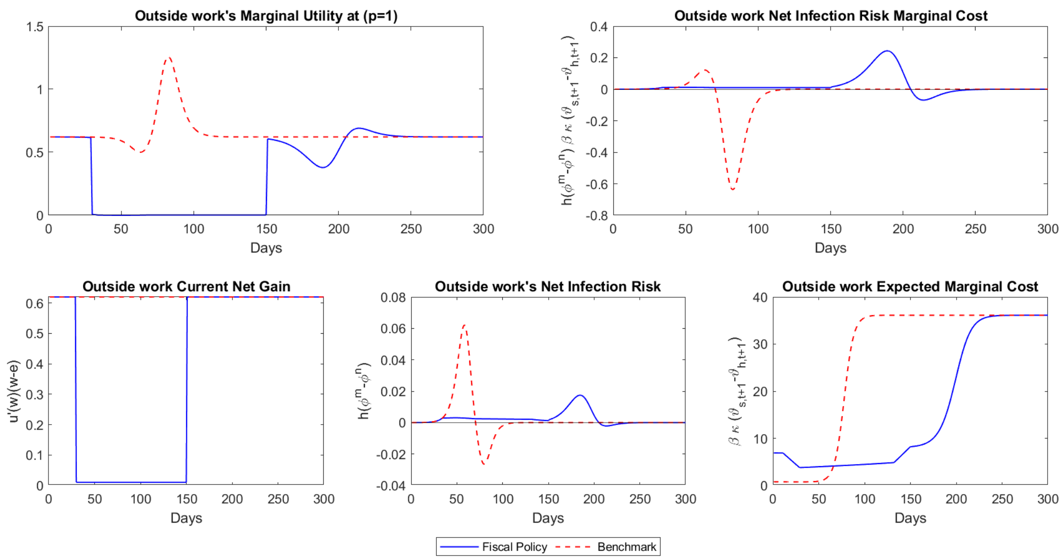

outside, given by $u'(w_t)(w_t - e_t)$. The fiscal policy reduced this incentive to virtually zero for four months. As a consequence, the marginal value of outside work fell during that period, as shown in the upper left panel. When the fiscal policy incentive expires, however, the value of outside work 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 fiscal policy 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 universal transfers to all agents in the population would essentially have 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, even manage to *flatten the curve*, but that impact is determined by the incentives embedded in the policy.

7 Final Remarks

In this paper we have explored a natural and intuitive idea: how fast a virus spreads across people depends on how they behave in response not only to the health costs of contagion but also to economic and financial incentives. The 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, for the scope for public health directives and economic policy to affect both lives and livelihoods during a pandemic.

In building the model in this paper we have made every effort to rely on first principles. As mentioned, the goal is that the model can be used for a reliable policy evaluation. Importantly, the model can be easily calibrated to economies with different structural characteristics, so as to provide specific policy evaluations.

An additional payoff of building a model from first principles is that the model can be modified in straightforward ways to tackle many additional important issues. Next we comment on the ones that appear most interesting or important to us.

We have abstracted from studying the effects of vaccination, but it should not be too hard to extend our model in that direction. For example, one can imagine that a vaccine becomes available 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. In this way, the expectation of vaccine development may have played a role in altering the duration of the Covid-19 episode. We believe that this is a promising question for future research.[14]

We have followed recent literature and convention in economics and assumed that agents have perfect foresight. One could argue that 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 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. On the other hand, the Covid-19 episode was not the first pandemic. Previous experiences, for example that of the Spanish Flu a century prior, were useful to anticipate many aspects of the evolution of the Covid-19 virus. Given this, our position is that assuming perfect foresight is a reasonable starting point for a model like ours. For further research, one might ask how much the analysis might change when the perfect foresight assumption is replaced, for instance, by a learning process.

Of the many possible extensions of our model, an interesting one would be to develop the implications of heterogeneity. For example, allowing inequalities over the $w_t - e_t$ wedge 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 provide 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.

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

References

Acemoglu, D., Chernozhukov, V., Werning, I., and Whinston, M. D. (2021). Optimal targeted lockdowns in a multigroup sir model. *American Economic Review: Insights*, 3(4):487–502.

Alvarez, F., Argente, D., and Lippi, F. (2021). A simple planning problem for covid-19 lock-down, testing, and tracing. *American Economic Review: Insights*, 3(3):367–382.

Brodeur, A., Gray, D., Islam, A., and Bhuiyan, S. (2021). A literature review of the economics of covid-19. *Journal of economic surveys*, 35(4):1007–1044.

Chang, R. and Velasco, A. (2020). Economic policy incentives to preserve lives and livelihoods. Technical report, National Bureau of Economic Research.

Chen, F. (2012). A mathematical analysis of public avoidance behavior during epidemics using game theory. *Journal of theoretical biology*, 302:18–28.

Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2021). The macroeconomics of epidemics. *The Review of Financial Studies*, 34(11):5149–5187.

Ganong, P., Noel, P., and Vavra, J. (2020). Us unemployment insurance replacement rates during the pandemic. *Journal of public economics*, 191:104273.

Garibaldi, P., Moen, E. R., Pissarides, C. A., et al. (2020). Modelling contacts and transitions in the sir epidemic model. *Covid Economics*, 5:1–21.

He, D., Zhao, S., Lin, Q., Zhuang, Z., Cao, P., Wang, M., and Yang, L. (2020). Therelative transmissibility of asymptomatic cases among close contacts. *International Journal of Infectious Diseases*, 1016.

Jones, C., Philippon, T., and Venkateswaran, V. (2021). Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies*, 34(11):5188–5223.

Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721.

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

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

Lucas, R. E. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, volume 1, pages 19–46. North-Holland.

Mossong, J., Hens, N., Jit, M., Beutels, P., Auranen, K., Mikolajczyk, R., Massari, M., Salmaso, S., Tomba, G. S., Wallinga, J., et al. (2008). Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS medicine*, 5(3):e74.

Rachel, L. (2022). An analytical model of behavior and policy in an epidemic. *American Economic Journal: Macroeconomics*.

Toxvaerd, F. (2020). Equilibrium social distancing. *Covid Economics*, page 110.

Verity, R., Okell, L. C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N., Cuomo-Dannenburg, G., Thompson, H., Walker, P. G., Fu, H., et al. (2020). Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet infectious diseases*, 20(6):669–677.

Weiss, H. H. (2013). The sir model and the foundations of public health. *Materials matematics*, pages 0001–17.