

DISCUSSION PAPER SERIES

DP14614

ECONOMIC POLICY INCENTIVES TO PRESERVE LIVES AND LIVELIHOODS

Andrés Velasco and Roberto Chang

**INTERNATIONAL MACROECONOMICS AND FINANCE
PUBLIC ECONOMICS**



ECONOMIC POLICY INCENTIVES TO PRESERVE LIVES AND LIVELIHOODS

Andrés Velasco and Roberto Chang

Discussion Paper DP14614

Published 14 April 2020

Submitted 13 April 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- International Macroeconomics and Finance
- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Andrés Velasco and Roberto Chang

ECONOMIC POLICY INCENTIVES TO PRESERVE LIVES AND LIVELIHOODS

Abstract

The Covid-19 pandemic has motivated a myriad of studies and proposals on how economic policy should respond to this colossal shock. But in this debate it is seldom recognized that the health shock is not entirely exogenous. Its magnitude and dynamics themselves depend on economic policies, and the explicit or implicit incentives those policies provide. To illuminate the feedback loops between medical and economic factors we develop a minimal economic model of pandemics. In the model, as in reality, individual decisions to comply (or not) with virus-related public health directives depend on economic variables and incentives, which themselves respond to current economic policy and expectations of future policies. The analysis yields several practical lessons: because policies affect the speed of virus transmission via incentives, public health measures and economic policies can complement each other, reducing the cost of attaining desired social goals; expectations of expansionary macroeconomic policies during the recovery phase can help reduce the speed of infection, and hence the size of the health shock; the credibility of announced policies is key to rule out both self-fulfilling pessimistic expectations and time inconsistency problems. The analysis also yields a critique of the current use of SIR models for policy evaluation, in the spirit of Lucas (1983).

JEL Classification: N/A

Keywords: crisis, COVID-19, Pandemics

Andrés Velasco - a.velasco1@lse.ac.uk
London School of Economics and CEPR

Roberto Chang - chang@economics.rutgers.edu
Rutgers University

Acknowledgements

Work on this paper was carried out while Roberto Chang served as BP Centennial Professor at the London School of Economics and Political Science. We acknowledge with thanks very useful conversations on the subject of this paper with several LSE colleagues, with Ricardo Hausmann from the Harvard Kennedy School, and with Norman Loayza from the World Bank. All errors are our own. Emails: chang@econ.rutgers.edu; A.Velasco1@lse.ac.uk

Economic Policy Incentives to Preserve Lives and Livelihoods*

Roberto Chang
Rutgers University and NBER

Andrés Velasco
London School of Economics and CEPR

April 2020

Abstract

The Covid-19 pandemic has motivated a myriad of studies and proposals on how economic policy should respond to this colossal shock. But in this debate it is seldom recognized that the health shock is not entirely exogenous. Its magnitude and dynamics themselves depend on economic policies, and the explicit or implicit incentives those policies provide. To illuminate the feedback loops between medical and economic factors we develop a minimal economic model of pandemics. In the model, as in reality, individual decisions to comply (or not) with virus-related public health directives depend on economic variables and incentives, which themselves respond to current economic policy and expectations of future policies. The analysis yields several practical lessons: because policies affect the speed of virus transmission via incentives, public health measures and economic policies can complement each other, reducing the cost of attaining desired social goals; expectations of expansionary macroeconomic policies during the recovery phase can help reduce the speed of infection, and hence the size of the health shock; the credibility of announced policies is key to rule out both self-fulfilling pessimistic expectations and time inconsistency problems. The analysis also yields a critique of the current use of SIR models for policy evaluation, in the spirit of Lucas (1983).

* Work on this paper was carried out while Roberto Chang served as BP Centennial Professor at the London School of Economics and Political Science. We acknowledge with thanks very useful conversations on the subject of this paper with several LSE colleagues, with Ricardo Hausmann from the Harvard Kennedy School, and with Norman Loayza from the World Bank. All errors are our own. Emails: chang@econ.rutgers.edu; A.Velasco1@lse.ac.uk

1. Introduction

Imagine you are a middle-aged person living in a middle-class neighborhood in one of the world's great cities. It has been a month since the government confined you and your family to your flat, and you are getting anxious. You are not one of those privileged professionals who can do all of your work online. Instead, you run the kind of small business that requires face-to-face contact, all day long. The government postponed some of your tax payments and the bank gave you a bigger credit line. But nonetheless your cash reserves are running low. As is the patience of your employees, who send ever more frequent messages asking when they will be able to go back to work. They understand taking public transport to get to their jobs is risky, but staying home with the prospect of much-reduced incomes is looking riskier still. Your business could afford to remain closed for another month if you were certain the economy would spring back to normal at the end of that period, but nowadays ... who can be certain about anything?

Much has been written since Covid-19 hit about the stark choices governments face between preserving lives and preserving livelihoods. Much less has been said about the equally stark choices regular citizens face. Yet in the end, what citizens do could be at least as important as what governments do in determining how, when and at what cost we overcome the pandemic.

If those regular citizens live not in prosperous New York, London or Milan, but in Manila, Sao Paulo or Lagos, the choices they face will be particularly unappealing. Initial income levels matter. Going for two or three months with reduced or no paycheck may be feasible for well-off families in rich countries, but not for households in developing countries whose incomes hover barely above the survival line. In the developing world the prevalence of informal jobs in informal firms further hinders the policy response, since governments may be unable to identify and get emergency financial aid to the workers and firms that need it. In the absence of past tax returns and accounting statements, banks may not be able to lend in order to tide people over until the crisis ends.

And where governments have been inept, corrupt, or both, citizens may ignore their entreaties to stay locked down—or to return to work when the time comes. Even worse: because people's willingness to forego income today hinges crucially on their confidence they will enjoy restored incomes in the future, trust in government policies, and the credibility of government announcements of an eventual recovery, are absolutely crucial for fighting the pandemic. But in countries where governments have seldom delivered on past promises, why should citizens believe them now?

To make sense of all of these complex and possibly conflicting factors, we need an *economic* theory of pandemics. And what the world has at its disposal today, for the most part, are *epidemiological* theories of pandemics. The difference is not just academic. *Epidemiological* theories are backward-looking: people's past choices determine how many cases of infection there are today. By contrast, *economic* theories are forward-looking: people's choices today—including the decision to engage or not in risky behavior that could result in infection—depend crucially on what they expect the future will bring.

An economic theory of pandemics is also necessary for the proper design of what government should and should not do during a pandemic. That is because economic policies can not only alleviate the economic and social effects of disease, but also change the severity of the pandemic itself. They can do so by changing the incentives that people face when making choices that, explicitly or implicitly, determine their risk of infection.

If this is so, then the analysis of alternative policies should take into account their possible incentive effects and the resulting impact on the dynamics of disease. Badly designed economic policies can be at odds with lockdowns, social distancing and other public health measures. But, our analysis shows, thoughtful economic and public health policies can also reinforce each other in reducing the impact of the pandemic.

To illuminate the feedback loops between medical and economic factors, we develop a minimal economic model of pandemics. In the model, epidemiological dynamics are similar to those in the standard SIR (susceptible-infected-recovered population) models. But, in contrast with those models, here contagion dynamics are affected by economic choices about whether to work or stay at home, today and in the future. In spite of its simplicity, the model yields interesting and sometimes unexpected results.

Unsurprisingly, the decentralized equilibrium of our economy is inefficient, because an externality is at work: when deciding whether or not to stay at home, people do not take into account the impact of their choice in the relative numbers of healthy and infected people “out there” in the workplace, and therefore on the overall speed of disease transmission.

Less obviously, the externality means that people can behave in a manner that is too risk-averse relative to the social optimum. If many infected people are at work already, and there isn't enough testing to identify them and compel them to stay home, then having one more person go to work could in fact reduce the share of infected people in the workforce, and therefore cut back on the risk of infection. Since people do not internalize this effect, they choose to stay away from work even in circumstances when this is not socially desirable.

Multiple expectational equilibria can also occur. If one person expects others to behave in such a way as to reduce the risk of infection, then it can pay off to ignore lockdown provisions and go to work. An equilibrium follows in which no one stays home. Conversely, the expectation that others will stay home can make it attractive to stay home, and society ends up in a full —and fully voluntary— lockdown. These equilibria can be Pareto ranked. We show the economy need not land in the outcome a benevolent social planner would have chosen. Depending on parameter values, full lockdown and no lockdown at all can be equilibrium outcomes, even when neither is optimal.

Using this model we then turn to the effects of alternative economic and public health policies. We show that several economic policies can make a difference not only for economic payoffs but also for health outcomes. One such policy is paying people to stay home during the infection period.

Such a transfer can induce more people to stay home, reducing contagion. But, we show, not just any payment will do. The transfer has to be large enough to induce expectations that other people will also stay at home. If too small, the transfer by itself will not succeed in eliminating the equilibrium in which everyone goes to work. Yet the transfer policy can work if complemented by fines on people who break a government-mandated lockdown. This illustrates how economic and public health policies can complement one another. An implication is that economic policies that provide appropriate incentives can reduce the costs of lockdowns — something government will like to hear, since the productivity and fiscal costs of generalized lockdowns are huge.

Strikingly, expectations of policies to be enacted after the initial contagion phase is over can matter for the extent of contagion itself. Any policy that causes people to expect higher future wages or, more generally, higher economic returns to being healthy —and therefore able to work— can induce individuals to stay home during the contagion phase. This is a novel reason to support expansionary policies to be implemented once the pandemic has peaked: if people come to expect them, they will have more reason to avoid infection today.

The danger, on the other hand, is that if people are pessimistic about the future they will behave today in ways that increase the risk of infection —and as a result make that pessimism self-fulfilling. Another danger is time inconsistency: after the pandemic has peaked the policymaker may find it that the cost of honoring the promise of wage subsidies or fiscal expansion is too high, and may therefore renege on the earlier announcement. This suggests that only governments with credible leadership and a history of respecting promises will be able to generate the kind of expectations of future policy that can help contain the pandemic today.

Finally, we also show large-scale testing to be a promising policy. But the conclusion comes with a twist: because testing reduces the risk of going to work, governments will have to pay people more to persuade them to stay home. So testing may have an indirect fiscal cost, unacknowledged so far.

The paper is structured as follows. Section 2 sets up our basic economic model of pandemics. Section 3 discusses the individual decision of whether to stay at home or working, and hence of how much exposure to infection risk is tolerable. Section 4 characterizes the general equilibrium of the model, while Section 5 contains a discussion of welfare aspects.

We develop our policy analysis in sections 6 and 7. Section 8 speculates on possible extensions, offers conjectures, and suggests additional implications. In that section we also relate our analysis to other existing work. Section 9 concludes.

2. A model of epidemics and economic incentives

Consider a simple model of an economy that lasts two periods, $t = 0, 1$. One can think of period 0 as the initial contagion stage and of period 1 as the recovery phase.

There is a continuum of agents. Population is constant and its size normalized to one. There are two locations we call “home” and “work”. Each individual who goes to work in period t produces a quantity w_t of a single final good, so total output in this economy depends on the number of people who work outside their homes. Normally everyone would go to work, but these are not normal times.

At the beginning of time a fraction $1 - h_0$ of the population is infected with a virus. The rest are healthy. Assume further that a fraction q of the whole population is impossible to reach or test. As a result, in period 0 people in that group always go to out to work. Being drawn randomly from the whole population, qh_0 are healthy and $q(1 - h_0)$ are infected.

The remaining $1 - q$ people are available for testing.¹ Assume for simplicity that all are tested. Naturally, $(1 - q)h_0$ are revealed to be healthy and $(1 - q)(1 - h_0)$ are revealed to be infected. Those who learn they are ill are compelled to stay home and remain isolated. But each healthy person must decide whether to stay home or go to work. Call these people “decision-makers”.

A decision-maker’s choice is not trivial. If she stays home she has given earnings e_0 , also in units of the good. She earns w_0 if she goes to work, but can contract the virus if she meets an infected person. As with all infected agents, the decision-maker must stay home at $t = 1$ if she gets the virus. In addition to being unable able to work, infected people suffer a utility loss χ . This is meant to capture the direct pain and suffering associated with illness.

The key aspect of this model is that the evolution of contagion is determined by people’s choices, which are health choices but also economic choices. To see this, let p denote the fraction of decision-makers who choose to go to work, and let be ϕ the probability that a healthy individual who goes to work does not get infected with the virus. It follows that the number of healthy people in the final period is

$$h_1 = h_0 - (1 - \phi)[qh_0 + p(1 - q)h_0]$$

The previous expression is just like the key equation in the famous SIR model of infectious transmission (Kermack and McKendrick, 1927), except that here ϕ depends on p . This apparently minor difference turns out to be crucial, since p (and therefore ϕ) are *endogenous*: they depend on the choices of decision-makers and reflect expectations about economic policy.

Period 1 is very simple. Health status becomes public information at the end of period 0 and infection lasts until the end of period 1. So in that period $1 - h_1$ people are ill and must stay home, in which case they earn some amount e_1 . The remaining h_1 people are healthy and, assuming $w_1 > e_1$, they choose to work and earn w_1 .

¹ Here and in the remainder of the paper we refer to antigen testing —that is, testing to detect if a person is currently infected. There is also antibody testing, which detects whether a person has developed immunity to the disease.

How do people become infected? Healthy individuals can only catch the virus if they go to work in period 0. There is no contagion at home, reflecting the assumption that sick people are isolated.

In the workplace, the basic assumptions of the SIR model apply. When at work a person can be the victim of contagion by randomly meeting an already-infected co-worker. It follows that the probability that a healthy person at work is still healthy in the final period is simply equal to the percentage of the working population that is healthy. That is,

$$\phi = \frac{qh_0 + (1-q)ph_0}{q + (1-q)ph_0} < 1$$

The probability that a healthy agent who goes to work gets infected is then $1 - \phi$.²

With this definition the transition equation can be written as

$$h_1 = h_0 - \phi q(1 - h_0)$$

This is a little model of a huge phenomenon. Yet the model clearly illustrates a crucial interaction that has been virtually ignored in the literature: the dynamics of contagion depend, at least in part, on people's choices; and those choices depend on economic considerations —not only about current conditions, but also about future economic policies and outcomes.

3. Individual decisions

Consider the choice of a decision-maker in period 0. Staying home means that she receives earnings e_0 . In addition, since there is no contagion at home, she will be able to work in period 1 and earn the reward w_1 . Assume, for simplicity, that agents have linear utility and there is no discounting. Then the value to the decision-maker of staying at home is $e_0 + w_1$.

Alternatively, the decision-maker can go to work in period 0. But then she runs the risk of infection and being unable to work in period 1. Recalling that the probability of infection at work is $(1 - \phi)$, and that infection also causes a utility loss χ , the value of going to work in period 0 is $w_0 + \phi w_1 + (1 - \phi)(e_1 - \chi)$.

² A more general formulation would allow each individual at work randomly to meet $(1 + \rho)$ others also at work. In turn, a healthy person at work would get infected with some probability κ if she met an infected worker. To alleviate notation, we shall impose $\rho = 0$ and $\kappa = 1$ in the subsequent discussion, although it will be clear that more general cases are easy to analyze.

Hence a decision-maker will go work today if $w_0 + \phi w_1 + (1 - \phi)(e_1 - \chi)$ is greater than $e_0 + w_1$ or, equivalently, if

$$w_0 - e_0 > (1 - \phi)(w_1 - e_1 + \chi)$$

She will stay home if the opposite is true.

This inequality is crucial and establishes how individual agents choose their exposure to contagion depending on economic variables. Indeed, the inequality compares the current gain from going to work to the expected loss, the latter given by the probability of infection times the sum of two elements: foregone income and disutility in the case of infection.

A decision-maker's choice depends on a "double relative": today's value of working relative to staying at home, and tomorrow's overall welfare relative to today's. In other words, both intra-temporal and intertemporal considerations matter. This will be particularly important for our policy analysis.

Observe also that the decision-maker's choice depends on $(1 - \phi)$, the probability of infection at work. But, as we saw, that probability depends on how many decision-makers go to work. The final outcome is then determined by the equilibrium choices of all decision-makers.

4. Equilibria

An equilibrium is defined in the usual way. Given the linearity of the model, it is natural to start by asking whether there are equilibria with either $p = 0$ or $p = 1$.

Consider $p = 0$ first. In that case, $\phi = h_0$, reflecting the fact that if no decision-maker goes to work the probability of meeting a healthy agent in a random meeting is just equal to the frequency of the healthy in the initial working population. For this to be an equilibrium, a typical decision-maker must find it optimal to stay at home, which requires

$$w_0 - e_0 < (1 - \phi)(w_1 - e_1 + \chi)$$

or, with $\phi = h_0$,

$$\frac{w_0 - e_0}{w_1 - e_1 + \chi} < 1 - h_0$$

Since the term on the RHS is always less than 1, an equilibrium in which all decision-makers stay at home exists in this model under some parameter conditions.

What are those conditions? Observe first that the LHS is smaller, and the preceding inequality less restrictive, if χ is larger. This is only natural: if working outside the home can result in infection, decision-makers will choose to stay home if χ is sufficiently large.

Other factors are economic. The inequality is more likely to be satisfied if the relative cost of staying home today is small compared to the relative cost of staying home tomorrow. This is why the ratio on the LHS of the inequality increases with $w_0 - e_0$ and falls with $w_1 - e_1$.

Finally, the inequality is less restrictive if h_0 is small. In that case, the probability meeting an infected person at is large.

In an equilibrium with $p = 0$, the final number of infections is minimal. The transition equation for the share of healthy people becomes:

$$h_1 = h_0 - qh_0(1 - h_0)$$

Note that this equation again has the SIR form, but now with $\phi = h_0$.³

While infections are lowest in an equilibrium with $p = 0$, the implications for production are ambiguous. In the first period the number of people at work is as small as it can be, so total output is minimized in that period. On the other hand, the number of available workers and, hence, total output in the second period are both maximized. We elaborate on the policy implications of this tradeoff in a later section.

Can there be an equilibrium with $p = 1$? Analogous reasoning leads to the conclusion that the answer is yes if

$$\frac{w_0 - e_0}{w_1 - e_1 + \chi} > 1 - \frac{h_0}{h_0 + q(1 - h_0)}$$

The intuition is analogous as that of the case $p = 0$. But there is a key difference. In an equilibrium with $p = 1$ each decision-maker's perception of the probability of infection if she goes to work is different than in an equilibrium with $p = 0$. This is because of the "contagion technology": the proportions of healthy versus and workers depend on how many decision-makers go to work.

The period-1 shares numbers of infected and healthy people are again given by equations of the SIR type. The healthy evolve according to $h_1 = h_0 - \phi q(1 - h_0)$, but in this case

$$\phi = \frac{h_0}{q + (1 - q)h_0}$$

³ In this equilibrium, advocates of the SIR model would claim that the model would have been "right" if only they had been able to pin down the correct ϕ from past information. This is reminiscent of Lucas (1983), and also consistent with the current debate of shifting U.S. predictions for the impact of Coronavirus. See sections 8 and 9 for further discussion.

which is smaller than the ϕ in the $p = 0$ equilibrium. This underscores the fact that the dynamics of contagion depend on economic factors. One of the important economic determinants is people's expectations about the future, which play no role in SIR-type models.

In fact, multiple self-fulfilling expectational equilibria can exist in our model. Given the analysis above, equilibria with $p = 0$ and $p = 1$ are feasible provided that

$$(1 - h_0) \left(\frac{q}{q + h_0(1 - q)} \right) < \frac{w_0 - e_0}{w_1 - e_1 + \chi} < (1 - h_0)$$

The intuition is as follows. More people get infected with the virus if more decision-makers go to work instead of staying home. But if more decision-makers go to work and infection rates increase, the relative rewards of working relative to staying home change in favor of the former, inducing more decision-makers to work, even after taking into account the higher risk of getting sick. Conversely, if more people stay home, infection rates fall, and future economic conditions change so as to induce decision-makers to remain home. So there is strategic complementarity across decision-makers' actions, and that can produce multiple equilibria.

We do not want to overemphasize here the possibility of those multiple expectational equilibria. But we do wish to underscore the crucial role of expectations in determining the dynamics of contagion. The speed of contagion is not only a public health issue: it is an economic issue as well.

The main and key implication so far is that there is a two way interaction between public health outcomes (and, as we will see, policies) and economic variables, including policies. That the interaction goes both ways turns out to be crucial to think about policy. But so far in the discussion the link has only been recognized in one direction: economists have largely taken the dynamics of the epidemics as an exogenous shock, and tried to tweak the policy response to attenuate social costs. Our analysis reveals that feedback in the other direction can also matter: economic policies influence the severity of the pandemic itself.

We expand on the policy analysis shortly. But before it is necessary to take a stand on what is socially desirable in this situation. We now turn to that issue.

5. Welfare implications

Suppose that the social welfare function is nothing but the weighted average of individual welfare levels, with the weights provided by the shares of the population at work and at home. Therefore,

$$W = [q + p(1 - q)h_0]w_0 + [(1 - q) - p(1 - q)h_0]e_0 + h_1w_1 - (1 - h_1)\chi$$

Clearly, social welfare is a function of p , the share of people who go to work. Ask next which is the setting of p a benevolent planner would choose in order to maximize social welfare.

Appendix 1 shows that the sign of the derivative of W with respect to p is the sign of

$$(w_0 - e_0) - (w_1 + \chi - e_1)(1 - \phi)^2$$

So if $(w_1 + \chi - e_1)$ is sufficiently large relative to $(w_0 - e_0)$, then social welfare is always decreasing in p . This is because having more people go to work has two effects that point in the same direction: it increases the human cost of infection and also cuts back on the number of people healthy who can go to work in the future, when $(w_1 - e_1)$, the gain from working relatively to staying at home, is large. In this case, there is no tradeoff between protecting lives and protecting livelihoods: keeping at home everyone who can be compelled to do so is clearly the better policy.

Conversely, if $(w_1 + \chi - e_1)$ is small relative to $(w_0 - e_0)$, then staying at home has benefits but also costs, because it means foregoing the relatively large reward from working in period 0. In this case there is indeed a tension between protecting lives and protecting livelihoods.

Appendix 1 also shows that the second derivative of W with respect to p is always positive, meaning the social welfare function is convex, not concave, in the share of people to go to work. So there is no interior optimum. The choice for the benevolent planner is simply between $p = 0$ (everyone stays home) or $p = 1$ (everyone who can goes to work).

Which one is better? Appendix 2 shows that $p = 0$ is preferred if

$$\frac{w_0 - e_0}{w_1 - e_1 + \chi} < \gamma(1 - h_0)$$

where

$$\gamma = \frac{q(1 - h_0)}{h_0 + q(1 - h_0)} < 1$$

By contrast, under decentralized decision-making the condition was

$$\frac{w_0 - e_0}{w_1 - e_1 + \chi} < 1 - h_0$$

Two conclusions strike the eye. The first is that the condition for the planner is not the same as for the individual. That is not surprising, given that there is an obvious externality: in deciding to go or not to work, individuals do not take into account the impact their own decisions have on the aggregate infection risk.

The second conclusion is more surprising: because the coefficient γ is smaller than one, the condition for $p = 0$ to be optimal is more stringent for the planner than for the individual! Put differently, $(w_1 + \chi - e_1)$ has to be larger relative to $(w_0 - e_0)$ in the case of the planner. So individuals are more “conservative” (more inclined to stay home) than is socially optimal.

Why is that so? Because the proportion of infected people is larger among decision-makers than among people “out there” in the workplace, and what decision-makers fail to internalize is that if they go to work, they actually help reduce —not enhance— the risk of contagion at work.

The practical implication is that society could end up locked down even in situations in which that is not socially desirable to do so. That may seem far-fetched but isn't. Think of the UK, which initially tried to adopt a soft lockdown like the one Sweden has adopted, but soon gave up because of political pressure to “do more”. Or think of Chile, where mayors are constantly pressuring the national government to impose a more stringent lockdown, in more regions of the country, than the government thinks is necessary or desirable.

The point may also be relevant when the time comes to lift lockdown policies. So far the focus has been on persuading people to stay home. But eventually governments will also have to persuade people to go back to work. The analysis here suggests that the second task may turn out to be anything but easy or straightforward.

We have postulated that the social welfare function is the weighted average of the ex post utility of the individual agents in the economy, which turns out to coincide with the expected welfare *ex ante* (that is, prior to the contagion period) of the representative individual. One can argue, however, that social welfare can be different from expected individual welfare for a number of reasons. So the social cost of infection could be larger than the individual cost. One way to capture this possibility is to replace the parameter χ in the previous social function W by some $\hat{\chi} > \chi$, with the gap $\hat{\chi} - \chi$ capturing the discrepancy between the social and individual cost of infection.

With this change the analysis in this section remains the same except for one observation: if the social cost $\hat{\chi}$ is high enough, the social optimum entails minimizing the extent of contagion — that is, setting $p = 0$ no matter what. In the current debate it is often claimed that policy should aim to minimize the number of infections regardless of economic costs. Assuming a sufficiently large $\hat{\chi}$ is one way to formalize and justify that belief.

6. Policies during the contagion phase

Our model is extremely simple, but precisely because of that simplicity it helps identify and understand the implications of alternative policies —both of public health policies such as lockdowns and economic policies such as taxes and transfers.

Government policies require resources, which in turn have an alternative social value. To say something about desirable policies, one must take this alternative social value of resources into account. We do that in the following way.

Suppose the government that can impose taxes, make transfers and enact laws. Assume it can also provide a public good, which is purchased at the end of period 1 and has social utility value proportional to the amount spent. One can think of the public good as infrastructure, or international reserves, or the assets held in a sovereign wealth fund. Regardless of the precise interpretation, suppose that if the government invests i in the public good, all agents receive a utility bonus αi , where α is a positive constant. The assumption that the marginal value of the public good is constant keeps the analysis that follows manageable.

As for public finance, assume that at $t = 0$ the government has an initial endowment $f > 0$. Aside from transfers to households, the government has no other expenditures. It can borrow or lend at a zero interest rate (this is consistent with equilibrium because our assumptions ensure that private agents would in fact be willing to borrow and lend at a zero rate).

Imagine that in the absence of the virus the government would have imposed no taxes nor made any transfers. In that case, the size of the public good provided would be given by the size of the initial government reserve, so $i = f$. Utility for each person would be $w_0 + w_1 + \alpha f$.

Now consider what happens when the virus hits. As a benchmark, suppose that

$$\frac{w_0}{w_1 + \chi} > 1 - h_0$$

We showed earlier that under these parameter values the only equilibrium in the absence of government action is $p = 1$ and therefore the highest possible rate of infection. In that case the virus causes expected utility to fall to

$$[q(1 - h_0) + h_0]w_0 + h_1w_1 - \chi(1 - h_1) + \alpha f$$

with

$$h_1 = h_0 - \frac{qh_0(1 - h_0)}{q + (1 - q)h_0}$$

The absence of government action implies two kinds of losses: fewer agents work in both periods and in the end there are $(1 - h_1)$ infections, which inflict a direct utility loss $\chi(1 - h_1)$.

Now suppose instead that in response to the virus the government gives a transfer to people who stay at home in period 0. This means that the government makes e_0 positive instead of zero. What is the impact? With a positive e_0 , one might guess that expected utility would become

$$[q(1 - h_0) + h_0]w_0 + (1 - q)(1 - h_0)e_0 + h_1w_1 - \chi(1 - h_1) + \alpha[f - (1 - q)(1 - h_0)e_0]$$

with h_1 defined as in the previous equation. This is a natural conjecture. With $p = 1$, a group of people of size $(1 - q)(1 - h_0)$ would receive the transfer, which would have to be financed with an equivalent-size reduction in public good provision.

Would this policy be welfare-improving? The expression above reveals it would be if and only if α , the marginal value of the public good, is less than one. When $\alpha > 1$ a positive e_0 would not be justifiable, no matter how much direct pain and suffering the virus causes. Remarkably, this conclusion would follow independently of the values of w_0 , w_1 and χ .

The preceding analysis takes the dynamics of infection as “shock” to be confronted by economic policy. That is precisely what makes it wrong. It fails to acknowledge that e_0 alters economic incentives for people to stay home and, in so doing, it can cut the severity of the infection shock.

In particular, suppose that e_0 is chosen so that

$$(1 - h_0) \left(\frac{q}{q + h_0(1 - q)} \right) > \frac{w_0 - e_0}{w_1 + \chi}$$

Then, in equilibrium, p must fall to zero: no decision-makers go to work and expected utility is

$$qw_0 + (1 - q)e_0 + h_1w_1 - \chi(1 - h_1) + \alpha[f - (1 - q)e_0],$$

where, in this case,

$$h_1 = h_0 - qh_0(1 - h_0)$$

This last equation indicates that the number of infections falls to its lowest possible level.

Therefore, the correct analysis differs from the previous “incorrect” one by taking into account the changes that e_0 induces on individual choices. It recognizes that a large enough e_0 causes decision-makers to stay home. Total output changes: it falls in period 0 because fewer people go to work, and increases in period 1, because more people are healthy then. Last but certainly not least, the increase in h_1 has the direct beneficial effect of reducing pain and suffering.

The total fiscal cost of transfers would be $(1 - q)h_0e_0$, with a direct utility impact of $(1 - \alpha)(1 - q)h_0e_0$. If $\alpha < 1$ there is no tradeoff. But if $\alpha > 1$ the policy has a direct utility cost, which has to be weighed against the indirect utility benefit caused by the change in behavior. The rewards for work w_0 and w_1 and the direct cost of infection χ are important to evaluate overall welfare effects. Indeed, for a large enough χ , it would be optimal to set a positive e_0 .

The general point is that economic policies can have incentive effects on the individual decision problems about how much infection risk to take, which can induce agents to change their choices in a way that alters the dynamics of infection. Designing economic policy to deal with a pandemic must take this possibility into account, for it can alter the relative evaluation of alternative policies. At the same time, the potential change in behavior can provide opportunities to reduce the human cost of the crisis if measures are properly tailored.

Consider, for example, the impact of cash transfers. In our model, a general cash transfer in period 0 is simply a gift of some size, say τ , to all agents. A moment's thought reveals that the transfer's impact on expected utility is just $(1 - \alpha)\tau$, the difference between the consumption value of the transfer to agents minus the value of the cost of the transfer (the reduction of the size of the public good i). What about the impact on equilibrium infection rates? There is none, since a general cash transfer does not have any incentive effects: it increases the payoff of deciders by τ regardless of whether they work or stay home.

In contrast, a policy of giving transfers to individuals that stay home in the initial period does have incentive effects, as we saw before. This is the traditional argument for targeted cash transfers, except that here "targeted" acquires a particular meaning: transfers are more effective if allocated to people who stay home, because in addition to compensating for lost income they reduce the relative reward of work and hence the risk of infection and the spread of the virus.

The same argument applies to expanded unemployment insurance benefits. If viewed simply as a way of propping up household consumption, the policy is imperfect in that it only reaches people who already had jobs. Our model reminds us that increased unemployment insurance benefits have an additional effect: they can induce workers to stay home. In fact, e_0 could be interpreted as unemployment insurance.

So far we have asked how government can induce people to stay home by means of economic incentives. But "lockdown" policies can also be non-voluntary. Suppose that the government enforces a lockdown by imposing a penalty π on anyone found working during the contagion period. We can think of π as a fine, jail time, or perhaps social sanction like public shame. In all these interpretations π , if paid in equilibrium, involves a deadweight loss.

A small π would not induce the low-infection equilibrium, but it is obvious that a large enough π would induce decision-makers to stay home, resulting in the low-infection equilibrium. If this equilibrium is preferred to the high-infection one, the lockdown has a beneficial effect.

But it can also result in a deadweight loss. If you interpret q in our model as the fraction of the population that must work during the contagion period (for instance, those workers perform essential functions), then the social deadweight loss of the lockdown is at least $q\pi$. In fact, it could be larger. If incentives to work are strong enough, then as many as $q + (1 - q)ph_0$ could choose to pay the fine, in which case the deadweight loss would be $[q + (1 - q)ph_0]\pi$.

This suggests that lockdowns enforced by fines can be a very blunt way to slow the spread of the virus. But it also reveals a more subtle point: economic policies can make a lockdown more effective. To see this, suppose that the penalty π is too small in the sense that equilibrium would remain at $p = 1$. Clearly a sufficiently large increase in the period-1 transfer e_0 would ensure that equilibrium shifts to $p = 0$. Hence the transfer would render the otherwise counterproductive lockdown effective at reducing infection. And the reverse also holds: a transfer e_0 can be too small by itself to ensure the low infection equilibrium, but can attain that equilibrium if a lockdown policy is added. In other words, economic policy and health policy can be *complements*.

7. Contagion and expectations of policies during the recovery

People's choices in period 0, and hence the dynamics of infections, depend not only by policies on that same period. They also depend on expectations about policies and economic outcomes in period 1. This is so because decision-makers face an intertemporal trade-off: working today, and earning more income than if staying home, puts them at risk of infection and not being able to work tomorrow. Hence their choice depends on policies that affect the relative reward to work tomorrow, and on w_1 , the expected size of that reward.

For instance, consider a transfer of size τ to people who work in period 1. If the transfer is large enough, in the sense that

$$(1 - h_0) \left(\frac{q}{q + h_0(1 - q)} \right) > \frac{w_0}{w_1 + \tau + \chi},$$

then the high-contagion, $p = 1$ equilibrium disappears and the only feasible outcome is the one with $p = 0$. Intuitively, expectations of a high reward to work in the final period increase the expected cost of infection, inducing the individual decision-maker to stay home.

We can think of the period 1 reward to work as resulting from expansionary macroeconomic policy in the recovery phase. For example, one could append to our simple economy a macroeconomic model. The government would have the ability to increase w_1 by means of a fiscal expansion presumably financed by borrowing, implying a final fall in the size of the public good i . In such a situation, a fiscal expansion would work exactly as the τ transfer just described.

If the low-infection equilibrium is socially preferred to the high-infection equilibrium, then there is a new argument in favor of expansionary policies during the recovery phase: such policies might be desirable not only because of their impact on output and wages during the infection phase, but also because they affect incentives and help lower the rate of infection, increasing the size of the workforce available during the recovery.

This argument is quite intuitive. People will be willing to be "locked down" and to forego income in the short run if they expect the initial sacrifice will have a payoff in terms of both reduced infection rates *and* higher chances for well-paid employment once the crisis begins to abate. But if people come to expect the economy will be in sorry shape and their own incomes will be low in the future, then it is quite possible they will feel compelled to go out and try to increase their earnings today. That, of course, will increase the rate of infection, lower h_0 and reduce aggregate income tomorrow, thus ensuring the economy will indeed be in bad shape!

This line of reasoning suggests a second avenue for strategic complementarities and possible multiple equilibria in this model. So far we have treated w_1 , the wage in the recovery phase, as exogenous and independent of the number of people who are healthy and able to work in period 1. But it could well be that there are some economic activities that have fixed costs of operation, so

that they need minimum numbers of workers and/or customers to restart. In that case the average wage could be increasing in h_1 , the share of healthy people in period 1.

The rest of the story is easy to tell. If people expect a high monetary reward from being healthy and being able to work in the future, then they will be more likely to stay home and reduce the chance of infection today, enlarging the workforce tomorrow and making a buoyant economy possible. At the same time, the pickup in economic activity would increase government tax revenues, making it more feasible for the government to undertake expansionary policies and deliver another round of growth vitamins.

The lesson, then, is that people's willingness to behave cautiously during the emergency depends crucially on expectations of what will happen to the economy and what government policies will be during the recovery phase. But before we get too excited at the prospect of self-confirming cycles of optimism and infection abatement, one warning: promises of future expansion may be time-inconsistent, and therefore less than fully credible.

This is clear in the case of $\alpha > 1$. Then, when the recovery phase arrives the government will have an incentive to reduce the size of any transfer τ it may have previously promised: each unit saved in transfers would allow for a unit increase in the size of the public good, implying a net utility gain of $(\alpha - 1)$. Since at that point the final number of infected people has already been determined, a benevolent government would deliver smaller transfers than announced earlier. As we know full well from the literature on time inconsistency, government promises during the initial contagion phase would then be likely to be ignored, unless the government has some way to commit not to break those promises during the recovery phase.

So we conclude that policies expected for the recovery phase can have a significant influence on individual choices during the initial crisis phase and, consequently, on the dynamics of contagion and the spread of the virus. At the same time, for those policies to have an impact, they have to be announced in advance and be credible. This suggests that only governments with enough credibility, built via durable institutions or a history of honoring promises, can take advantage of this policy option and use it to diminish contagion and make progress against the pandemic.

8. Extensions and relation to existing work

How general are our results? At some level, they are very general: they come from marrying the standard account of infection transmission with a minimal model of economic behavior that responds to incentives. And they follow from a very simple but intuitive principle: if economic policies can affect those incentives, then they can affect the transmission of disease. So our approach to policy should apply in any model where the dynamics of infection depend, at least to some degree, on individual economic choices.

While our main messages do not depend on the particulars of the model (which is why we endeavored to convey them in the simplest model we could think of), it is worth thinking about the role of different features of the model and how some of our more specific results may depend on them.

Many of our assumptions, such as the linearity of preferences or a fixed output per worker, should not be too hard to relax and are unlikely to change anything substantial in the analysis.

More fruitful is to examine the implications of changing the technology of infection, which is the more novel part of our model. We started by assuming that a share q of the population cannot be reached or tested, while the rest of the population can be. As a result, so-called decision-makers are drawn from a set of people who know they have been infected. When those decision-makers choose to go to work, they enlarge the share of healthy people in the workforce. That is why ϕ , the probability of being infected at work, is increasing in p , the share of decision-makers who do work.

That specification is special in several ways. It implies that the share of decision-makers who are infected is less than the share of infected people “out there” in the workplace (or the street, the store, the public bus or the subway cart). That seems like a reasonable description of society in the early stages of spread of the virus, when many people are yet to be reached by authorities, informed of what is going on, and tested. But it also means that sending more decision-makers to work during the contagion period changes the proportions of healthy people versus infected ones in favor of the former. And while all decision-makers know themselves to be healthy, they do not internalize the impact of their individual choices on the aggregate probability of infection. This combination of features explains why there may be strategic complementarities, allowing for multiple equilibria and for the peculiar nature of externality effects.

We can ask, then, what happens if the parameter q falls. This is a particularly interesting question because a lower q is tantamount to an increase in the scale of testing, which Paul Romer and others have forcefully advocated.⁴ Leaving aside issues of feasibility and scalability of large-scale testing, what are the effects of that policy?

For any given p , in the model the effect of q on the probability that a decision-maker stays healthy is given by:

$$\frac{\partial \phi}{\partial q} = - \frac{h_0 p (1 - h_0)}{[q + (1 - q) p h_0]^2} < 0$$

So more testing means fewer chances of being infected and stronger incentives to go to work.

⁴ See, for example, Romer and Garber (2020).

Assume that the government attempts to induce decision-makers to stay home by giving them a transfer during the contagion period. Recall that the transfer would be always (that is, regardless of expectations) effective at attaining such objective only if

$$(1 - h_0) \left(\frac{q}{q + h_0(1 - q)} \right) > \frac{w_0 - e_0}{w_1 + \chi}$$

Clearly, the LHS is an increasing function of q .

Therefore, in the presence of more testing, the transfer necessary to keep people home is larger! This may seem surprising but is not: if it is now less risky to go to work, people will demand a larger compensation to stay home.

So large-scale testing is not without wrinkles. Or, to put it differently, to the direct logistical costs of testing millions of people every few days one should add the higher indirect cost of compensating them, once they have been tested, to make sure they do not go to work.

An extension would need to place the extent of testing at the center of our model. For instance, different shares could be tested within the q -population that must always go to work and the $(1 - q)$ population that can choose whether to do so. That could reverse the nature of the externality at work, and cause equilibrium behavior to be excessively risk-taking.

Observe also that we have assumed that, after testing, individuals not only learn if they are infected or not, but also forced to stay home if they turn out to be sick. In the parlance of Romer and others, we have assumed a “testing-cum-isolation” policy. But one could also ask what would happen if there is random testing but no isolation requirement. This may be a realistic alternative.

In our model analyzing a testing-without-isolation policy would require asking what would happen if infected individuals had the choice of working or staying at home in the initial period. Presumably the answer would depend on assumptions about how the rewards from working are affected by sickness—in a plausible case, w_t would be lower for infected people. This extension may be worth pursuing, but falls outside the main focus of our paper.

Likewise, we have assumed that individuals are *ex ante* identical. A straightforward extension of our model could allow for *ex ante* health heterogeneity. We assumed so far that at work a healthy person gets infected with probability one if she meets an infected coworker. Instead one could assume that the probability is given by some idiosyncratic parameter θ , where θ has a non-degenerate distribution in the population. This would capture the fact that people have different degrees of susceptibility to contagion, perhaps because of age.

An educated conjecture is that appropriate assumptions about the distribution of θ would introduce enough smoothness into the model so that equilibria other than the extreme $p = 0$ and $p = 1$ would appear. An interior equilibrium, with $0 < p < 1$, would have the appealing

property that only *ex-ante* healthier decision-makers (that is, those with θ greater than a certain threshold, if higher θ corresponds to lower probability of infection) would work.

Guessing other consequences of this extension is more hazardous. It is not obvious, for instance, that multiple equilibria would disappear; on the contrary, equilibria could become even more abundant. These questions offer promising avenues for future research.

Clearly many other extensions are imaginable, and one could endlessly conjure alternative formulations and their implications. To a large extent this is the case because the literature on the economics of pandemics is in its very early infancy. So rather than speculating further on variations of the model, it seems more fruitful at this point to clarify how our paper contributes to the related literature.

Our paper is most closely related to the very recent papers on macroeconomic policy responses to the Covid-19 crisis. Prominent examples are Fornaro and Wolf (2020), Faria e Castro (2020), and Jorda, Singh and Taylor (2020). These and other papers attempt to characterize the dynamics of the economy under alternative macroeconomic policies in an infinite-horizon setting. But they all take the dynamics of the Covid-19 pandemic as exogenously given.

Fornaro and Wolf (2020) argue that the pandemic can be seen as an adverse shock to productivity growth, while in Faria e Castro (2020) pandemics are tantamount to preference shocks. By contrast, our model has only two periods and is much less ambitious in terms of dynamics, but it does characterize the dependence of contagion on individual economic decisions. And, to the extent that individual decisions depend on current and future policy, the evolution of the pandemic can be influenced by policy. In this sense, modeling the impact of Covid-19 as exogenously given shocks potentially leads to invalid policy analyses. But the quantitative significance of this shortcoming remains, of course, an empirical issue.

Two recent papers include a channel for policy to influence virus dynamics through its impact on the decisions of individuals: Eichenbaum, Rebelo and Trabandt (2020); Jones, Phillipon, and Venkateswaran (2020).⁵ They all develop dynamic models where a virus hits the economy and contagion follows SIR-type dynamics. Moreover, the SIR equations are assumed to depend on economic activities such as consumption and hours worked. As in our analysis, individual agents understand that their consumption and labor supply choices have implications for their exposure to contagion.

But there are several differences between those three papers and ours. The main focus of Eichenbaum et al. and Jones et al., for instance, is on describing and quantifying dynamic implications, while our goal is to clarify and explore the channels through which policy can affect behavior and therefore the transmission of infection.

⁵ See also Kaplan, Moll and Violante (2020).

Perhaps most consequential for the study of policy are differences in the specific way economic activity affects contagion. Eichenbaum et al. and Jones et al. both assume that contagion increases with the levels of aggregate consumption and production. They do not provide a microeconomic justification for that assumption, but simply take it as a reduced form.

In contrast, our paper provides an explicit environment where SIR-type dynamics emerge endogenously. This difference turns out to matter for several results. For example, in Eichenbaum et al. and Jones et al., increasing consumption taxes during the contagion phase of a pandemic would reduce infections, which is an argument in favor of such a policy. In our model, consumption taxes have no impact on individual choice and therefore no effects on contagion dynamics.

The gravity of the Covid-19 epidemic has motivated a myriad of policy related studies and proposals. An influential collection is Baldwin and Weder di Mauro (2020). Loayza and Pennings (2020) provide a useful overview of policy issues, with emphasis on developing countries. Gourinchas (2020) discusses the need to coordinate economic responses and health policy, but he does not provide a formal analysis.

Macroeconomics *aficionados* will recognize a close connection between our analysis and the influential Lucas (1983) critique of econometric policy evaluation. For a given set of economic and health policies, any equilibrium of our model implies that infection dynamics are similar to those of the SIR model, whose parameters are a function of “deeper” underlying aspects of the environment, including the given policies. An implication is then that the SIR parameters must change if policies change. That implication is not simply a theoretical curiosity. Rather, it can be critical for policy analysis.

Lucas argued that macroeconomic policy based on reduced form econometric equations, especially the Phillips Curve, was unsound, because the parameters of those equations would shift as people changed behavior in response to policy. Strikingly, he showed that this would be the case even if the predictive power of the econometric equations, estimated from data from previous episodes, was strong: the equations would become unstable and the parameters change in undefined ways when new policies were implemented. Substitute “reduced form econometric equations” for “SIR equations” and the Lucas Critique applies with force to the current situation. This is one key takeaway from our paper.

9. Final remarks

The success or failure of public policies to fight the pandemic will depend on whether those policies induce socially-desirable patterns of behavior among ordinary citizens. And how people choose to behave in turn depends on a myriad of factors, including not only expectations of future policies, but also expectations of how other people will respond to those policies.

This paper provides a minimal model to understand the feedback loops involving economics, public health and expectations. One lesson is that it is easy for things to go frightfully wrong. If people come to be pessimistic—for instance, about the extent of contagion or the future health of the economy—they can react in ways that will make that pessimism self-fulfilling.

But another lesson is that there are policies that can potentially avoid those pitfalls. This paper has studied and characterized some of those policies in a minimal setting. Doing so in a richer environment, where the quantitative aspects of those policies can be explored more fully, is clearly a priority for future research.

We claim that economic policy can change the dynamics of contagion via its impact on incentives. How important is that link likely to be in practice? Several aspects of the current crisis suggest it can be quite important indeed.

Look at the massive change in recent projections of Covid-19 deaths in the United States. As recently as the end of March, the Trump administration was publicly projecting virus-related deaths between 100,000 and 240,000 by the end of the U.S. Summer. Less than two weeks later, the official estimate came down to just 60,000. This much-lower number is comparable to the usual number of deaths caused by influenza every year.

What explains the astonishingly large and sudden change in the official estimates? According to health officials and commentary by public health experts, the previous dire predictions assumed low compliance rates with lockdown and social distancing measures. That assumption turned out to be wrong: compliance by the U.S. population has been much better than expected, and that accounts for the bulk of the change in death estimates.⁶

This is, of course, excellent news. By now it should be obvious that whether or not people adhere to instructions to stay at home is crucial for the number of deaths from the virus. It is less obvious that those individual decisions are likely to be influenced by economic factors and policy incentives. In fighting the pandemic, policymakers and economists will be well served by remembering that fact and its implications.

⁶ This is an ongoing debate that has been widely reported by the press. See, for example, “US Coronavirus Predictions Have Shifted. Here is Why”, CNN.com, April 9 2020.

References

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

Eichenbaum, M., S. Rebelo and M. Trabandt, "The Macroeconomics of a Pandemic". NBER Working Paper 26882, Revised April 2020.

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.

Gourinchas, P.O., "Flattening the Curve of Infection and the Curve of Recession at the Same Time", *Foreign Affairs* (Web Edition), March 26 2020

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.

Kaplan, G., B. Moll and G. Violante, "Pandemics According to HANK". Powerpoint presentation, LSE, 31 March 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.

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.

Romer, P. and A. Garber, "Will Our Economy Die From Coronavirus?", *The New York Times*, March 23 2020.

Appendix 1

Welfare is

$$W = [q + p(1 - q)h_0]w_0 + [(1 - q) - p(1 - q)h_0]e_0 + h_1(w_1 + \chi) - (1 - h_0)\chi$$

Using the transition equation GDP can be written as

$$W = e_0 + e_1 + [q + p(1 - q)h_0](w_0 - e_0) + [h_0 - \phi q(1 - h_0)](w_1 + \chi - e_1) - (1 - h_0)\chi$$

It follows that

$$\frac{\partial W}{\partial p} = (1 - q)h_0(w_0 - e_0) - q(1 - h_0)(w_1 + \chi - e_1) \frac{\partial \phi}{\partial p}$$

which, substituting in the value of $\frac{\partial \phi}{\partial p}$, becomes

$$\frac{\partial y}{\partial p} = (1 - q)h_0[(w_0 - e_0) - (w_1 + \chi - e_1)(1 - \phi)^2]$$

So W can be either increasing or decreasing in p .

Notice also that

$$\frac{\partial^2 y}{\partial p^2} = 2(1 - q)h_0(w_1 + \chi - e_1)(1 - \phi) \frac{\partial \phi}{\partial p}$$

$$\frac{\partial^2 y}{\partial p^2} = \frac{2(w_1 + \chi - e_1)(1 - q)^2 h_0^2 (1 - \phi)^3}{q(1 - h_0)} > 0$$

So if $(w_1 + \chi - e_1)$ is sufficiently large relative to $(w_0 - e_0)$, W is a U-shaped function of p , with a minimum at

$$\phi = 1 - \left(\frac{w_0 - e_0}{w_1 + \chi - e_1} \right)^{\frac{1}{2}}$$

Finally, note that

$$y(p = 0) = qw_0 + h_0 [1 - q(1 - h_0)]w_1$$

$$y(p = 1) = [q + (1 - q)h_0]w_0 + h_0 \left[\frac{h_0}{h_0 + q(1 - h_0)} \right] w_1$$

Appendix 2

Welfare if $p = 0$ is

$$W(p = 0) = e_0 + e_1 + q(w_0 - e_0) + h_0 [1 - q(1 - h_0)](w_1 + \chi - e_1) - (1 - h_0)\chi$$

Welfare if $p = 1$ is

$$W(p = 1) = e_0 + e_1 + [q + (1 - q)h_0](w_0 - e_0) + h_0 \left[\frac{h_0}{h_0 + q(1 - h_0)} \right] (w_1 + \chi - e_1) - (1 - h_0)\chi$$

It follows that

$$\begin{aligned} W(p = 0) - W(p = 1) \\ = (1 - q)h_0 \left\{ (1 - h_0) \left[\frac{q(1 - h_0)}{h_0 + q(1 - h_0)} \right] (w_1 + \chi - e_1) - (w_0 - e_0) \right\} \end{aligned}$$

So from the point of view of the planner it is best to set $p = 0$ and keep everyone home if

$$\frac{w_0 - e_0}{w_1 - e_1 + \chi} < \gamma(1 - h_0)$$

where

$$\gamma = \frac{q(1 - h_0)}{h_0 + q(1 - h_0)} < 1$$