DISCUSSION PAPER SERIES

DP14606

HEALTH VERSUS WEALTH: ON THE DISTRIBUTIONAL EFFECTS OF CONTROLLING A PANDEMIC

Andrew Glover, Jonathan Heathcote, Dirk Krueger and Jose-Victor Rios-Rull

LABOUR ECONOMICS

MACROECONOMICS AND GROWTH

MONETARY ECONOMICS AND FLUCTUATIONS

PUBLIC ECONOMICS

ISSN 0265-8003

HEALTH VERSUS WEALTH: ON THE DISTRIBUTIONAL EFFECTS OF CONTROLLING A PANDEMIC

Andrew Glover, Jonathan Heathcote, Dirk Krueger and Jose-Victor Rios-Rull

Discussion Paper DP14606
Published 12 April 2020
Submitted 10 April 2020

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

HEALTH VERSUS WEALTH: ON THE DISTRIBUTIONAL EFFECTS OF CONTROLLING A PANDEMIC

Abstract

Many countries are shutting non-essential sectors of the economy to slow the spread of COVID-19. The gains and losses from these policies are very unequally distributed. Older individuals have most to gain from slowing virus diffusion. Younger workers in sectors that are shuttered have the most to lose. In this paper we first extend a standard epidemiological model of disease progression to include heterogeneity by age, and multiple sources of disease transmission. We then incorporate the epidemiological block into a multi-sector economic model in which workers differ by sector (basic and luxury) as well as by health status. Individuals value consumption, life, and health. We study optimal mitigation policies of a utilitarian government that can redistribute resources across individuals, but where such redistribution is costly. We show that optimal redistribution- and mitigation policies interact and thus the utilitarian government chooses a very different mitigation policy path than would be suggested by a representative agent setting. This policy reflects a compromise between the strongly diverging preferred policy paths across the subgroups of the population.

JEL Classification: N/A

Keywords: COVID-19, Economic Policy, redistribution

Andrew Glover - andyecon@gmail.com
Federal Reserve Bank of Kansas City

Jonathan Heathcote - grandeskimlatte@gmail.com
Federal Reserve Bank of Minneapolis

Dirk Krueger - dkrueger@econ.upenn.edu
University of Pennsylvania and CEPR

Jose-Victor Rios-Rull - vr0j@econ.upenn.edu
University of Pennsylvania

Acknowledgements
Thanks to Josep Pijoan for discussions about the value of life in health and in sickness. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Banks of Kansas City and Minneapolis or the Federal Reserve System.

Health versus Wealth: On the Distributional Effects of Controlling a Pandemic*

Andrew Glover† Jonathan Heathcote‡ Dirk Krueger§ José-Víctor Ríos-Rull¶

April 11, 2020

First version, March 31 2020

Abstract

Many countries are shutting non-essential sectors of the economy to slow the spread of COVID-19. The gains and losses from these policies are very unequally distributed. Older individuals have most to gain from slowing virus diffusion. Younger workers in sectors that are shuttered have the most to lose. In this paper we first extend a standard epidemiological model of disease progression to include heterogeneity by age, and multiple sources of disease transmission. We then incorporate the epidemiological block into a multi-sector economic model in which workers differ by sector (basic and luxury) as well as by health status. Individuals value consumption, life, and health. We study optimal mitigation policies of a utilitarian government that can redistribute resources across individuals, but where such redistribution is costly. We show that optimal redistribution- and mitigation policies interact and thus the utilitarian government chooses a very different mitigation policy path than would be suggested by a representative agent setting. This policy reflects a compromise between the strongly diverging preferred policy paths across the subgroups of the population.

Keywords: COVID-19; Economic Policy; Redistribution

*Thanks to Josep Pijoan for discussions about the value of life in health and in sickness. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Banks of Kansas City and Minneapolis or the Federal Reserve System.

†Federal Reserve Bank of Kansas City; Andrew.glover@kc.frb.org

‡Federal Reserve Bank of Minneapolis and CEPR; heathcote@minneapolisfed.org

§University of Pennsylvania, NBER, and CEPR; dkrueger@econ.upenn.edu

¶University of Pennsylvania, UCL, CAERP, CEPR and NBER; vr0j@upenn.edu

1 Introduction

The central debate about the appropriate economic policy response to the global COVID-19 pandemic is about how aggressively to restrict economic activity in order to slow down the spread of the virus, and how quickly to lift these restrictions as the pandemic shows signs of subsiding. In this paper, we argue that one reason people disagree about the appropriate policy is that "lock-down" policies have very large distributional implications. These distributional effects mean that different groups prefer very different policies. Standard epidemiological models assume a representative agent structure, in which households face a common trade-off between restrictions on social interaction that slow the virus transmission but which also depress economic activity. In practice, however, the benefits of slower viral transmission are not shared uniformly, but accrue disproportionately to older households, who face a much higher risk of serious illness or death conditional on becoming infected. At the same time, the costs of reduced economic activity are disproportionately born by younger households, who bear the brunt of lower employment. A second very important dimension of heterogeneity is between younger workers employed in different sectors. Sensible lock-down policies designed to reduce viral spread will naturally focus on reducing activity in sectors in which there there is a social aspect to consumption and sectors that produce goods or services perceived to be luxuries. For example, restaurants and bars are likely to be the first to be closed. Because workers cannot easily reallocate across sectors, this implies that lock-down policies will imply extensive redistribution between young households specialized in different sectors. Thus, different groups in the economy (old versus young, workers in different sectors, healthy versus sick) will likely have very different views about the optimal mitigation strategy. Furthermore, there is the potentially large need for redistributive public policies, and to the extent that these are costly to implement[1], the optimal macro mitigation policy will not only shape the need for, but in turn depend on redistribution policies at the micro level.

In this paper we seek to build and then quantitatively implement a framework to model this interaction between macro mitigation- and micro redistribution policies. This requires a structure with i) a household sector with heterogeneous individuals, ii) an epidemiological block that determines their health transitions through a potential epidemic, and iii) a government with tools for mitigation and redistribution, as well as a desire for social insurance.

[1]For example, the revenue for transfer programs needs to raised through distortionary taxation.

On the household side, we distinguish between three types of people: young workers in a basic sector, young workers in a luxury sector, and old retired people. Output of workers in the two sectors is combined to produce a single final consumption good. Workers are immobile across sectors, but sectoral output of the two sectors are perfectly substitutable. The only difference between the two sectors is that economic activity in the luxury sector impacts the rate of infections, which gives the policy maker a potential incentive to shut down part of the economic activity in that sector.

The epidemiological model builds on a standard SIR diffusion framework. We label our variant a *SAFER* model, reflecting the progression of individuals through a sequence of possible health states. Model individuals start out as susceptible S (i.e. healthy, but vulnerable to infection), can then become infected but asymptomatic A, infected with flu-like symptoms F, infected and needing emergency hospital care E, recovered R (healthy and immune), or dead. The three different types of model agent face differential infection risk (workers face more exposure than non-workers) and differential health outcomes conditional on infection (the old are more likely to end up in emergency care).

The government has a utilitarian social welfare function and two policy levers at its disposal to maximize social welfare. First, at each date the planner can choose what fraction of activity in the luxury sector to shut down. We call this policy the extent of mitigation. Mitigation slows the spread of the virus (by reducing the rate at which susceptible workers become asymptomatically infected), but it reduces income of workers in the luxury sector. Second, the planner chooses how much to redistribute income away from workers and toward those who are not working, either because they are old, because they are unwell, or because their workplaces have been closed due to mitigation. Redistribution is desirable due to the utilitarian social welfare function, but crucially, we also assume that this redistribution is costly, so that perfect insurance is not optimal. Conditional on a given path for mitigation, the optimal redistribution problem is equivalent to a static social planner problem, with lower aggregate consumption and more consumption inequality across workers the more costly is redistribution. This in turn feeds back adversely on the dynamic incentives for mitigation, implying that a government facing more costly redistribution needs will choose less mitigation dynamically.

In the context of the model with these trade-offs we compute optimal paths for mitigation, where (currently) the path for mitigation is restricted to be a simple parametric function of time. we find that the optimal path for mitigation is highly sensitive to the relative welfare

weights the government attaches to the three types of people in the model. A planner that prioritizes the old chooses extensive and prolonged mitigation, as the old are highly vulnerable to contracting the disease and dying from them. A planner that prioritizes workers in the luxury sector subject to a potential lock down chooses much weaker mitigation as the economic costs of foregone income and thus consumption dominate for this group.

We also consider how the optimal policy for a utilitarian equal-weights planner varies with the cost of redistribution across worker types. We find that the larger is this cost, the more moderate is optimal mitigation. Thus our economy which features key dimensions of inequality implies a more modest shutdown than a representative agent analogue, at the cost of higher mortality during the epidemic.

There is an extraordinary set of papers being written about the pandemic. To cite the ones that we are aware as of now: Atkeson (2020) was perhaps the first to introduce economists to the epidemiology SIR (Susceptible-Infectious-Recovered) class of models. He emphasizes the negative outcomes if and when the fraction of active infections in the population exceeds 1% (at which point the health system is forecast to be severely challenged) and 10% (which may result in severe staffing shortages for key financial and economic infrastructure) as well as the cumulative burden of the disease over an 18 month horizon. Greenstone and Nigam (2020) use the state of the art Imperial College epidemiological model (Flaxman et al. (2020)) to compare the paths under moderate distancing versus taking no action, and using the statistical value of life to assess the social cost of no action. They calculate 1.7 million lives saved between March 1 and October 1, 37% of them due to the avoidance of overcrowding in hospitals. They impute a benefit of such action of 58% of per capita yearly consumption.

Eichenbaum et al. (2020) extend the canonical epidemiology model to study the interaction between economic decisions and pandemics. They emphasize how equilibria without interventions lead to sub-optimally severe pandemics because infected people do not fully internalize the effect of their economic decisions on the spread of the virus. Krueger et al. (2020) argue that the severity of the economic crisis in Eichenbaum et al. (2020) is largely averted if individuals can endogenously adjust the sectors in which they consume. Toxvaerd (2020) characterizes the simultaneous determination of infection and social distancing. Moll et al. (2020) pose a version of a HANK model in which agents differ by occupation, and occupations have two key characteristics: how social their consumption is, and how easily work in the occupation can be done at home. They tie demand for social goods and willingness to work in the workplace to fear

of contracting the virus, with endogenous feedback to relative earnings by occupation. Bayer and Kuhn (2020) explore how differences in living arrangements of generations within families contribute to the cross country differences in terms of case-fatality rates. They document a strong positive correlation between this variable and the share of working-age families living with their parents. Berger et al. (2020) extend the baseline Susceptible-Exposed-Infectious-Recovered (SEIR) infectious disease epidemiology model to explore the role of testing and to thereby get a better idea of how to implement selective social separation policies. Fang et al. (2020) quantify the causal impact of human mobility restrictions using the Chinese experience and find that the lock-down was very effective, providing estimates of diffusion under different scenarios. Hall et al. (2020) provide a simple calculation to assess how much would we be willing to pay to have never had the virus (their answer is about a quarter of one year's worth of consumption).

We start by describing how we model the joint evolution of the economy and the population in Section 2. We then turn to describe the economic policies we have in mind in this environment (something that goes beyond social distancing) and how to think of the best policy Section 3. We then move to describe how we specify the quantitative details of our environment in Section 4.

2 The Model

We start describing the demographic structure in our *SAFER* extension of the standard *SIR* epidemiological model in Section 2.1. We then embed this component into a multi-sector production economy in section Section 2.2, describing how *mitigation* shapes the pattern of production, and how redistribution shapes the pattern of consumption. We defer the description of government mitigation and redistribution policies to Section 3.

2.1 Household Heterogeneity and Health Transitions in the SAFER Model

Agents can be young or old, which we denote y and o respectively. We think of the young as aged below 65 and making up the vast majority of the population. For simplicity, and given the short time horizon of interest, we abstract from population growth and ignore aging, i.e. the transition of young individuals into turning old.

Within each age group, agents are differentiated by health status that can take six different

values: susceptible s, asymptomatic a, miserable with flu-symptoms f, requiring emergency care e, recovered r, or dead d. Individuals in the first group have no immunity and are susceptible to infection. The a, f, and e groups all carry the virus and can pass it onto others. However, they differ in terms of their symptoms. The asymptomatic have no symptoms or very mild ones, and thus spread the virus unknowingly. We model this state explicitly (in contrast to the standard SIR type models) because a significant percentage of individuals infected with COVID-19 experience no symptoms.[2] Those with flu-like symptoms are sufficiently sick to know they are likely contagious and they stay at home and avoid the workplace. Those requiring emergency care are hospitalized. The recovered are again healthy, no longer contagious, and are immune from future infection. A worst case virus progression is from susceptible to asymptomatic to flu to emergency care to dead.[3] However recovery is possible from the asymptomatic, flu and emergency-care states. We model three sources of possible virus contagion: people can catch the virus from colleagues at work, from family or friends outside work, and from taking care of the sick in hospitals.

Young agents in the model are further differentiated by the sector in which they can work. A fraction of the young work in the basic sector, denoted b, while the rest work in a luxury sector, denoted ℓ. We assume that output of the basic sector is so vital that it is never optimal to barr even a subset of b sector workers from working. In contrast, it may be optimal to require some or all of the workers in the ℓ sector to stay at home in order to reduce the transmission of the virus in the workplace.[4] We will call such a policy a (macroeconomic) mitigation policy, m. More precisely, $m(t)$ will denote the fraction of luxury workers that are instructed to stay at home and not go to work at time t. We assume that workers cannot change sectors (at least not during the short time horizon studied in this paper); thus the sector of work is a fixed characteristic of a young individual.

Time starts at $t = t_0$ and evolves continuously. All economic variables, represented by roman letters, are understood to be functions of time, but we suppress that dependence whenever there

[2]DeCODE, a subsidiary of Amgen, randomly tested 9,000 individuals in Iceland. Of the tests that came back positive (1 percent), half reported experiencing no symptoms.

[3]Note that in the standard SEIR model, agents in the exposed state E have been exposed to the virus and may fall ill, but until they enter the infected state I they cannot pass the virus on. Our asymptomatic model state is a hybrid between the E and the I states in the SEIR model: asymptomatic agents have no symptoms (as in the SEIR E state) but can pass the virus on (as in SEIR the I state). Berger et al. (2020) make a similar modeling choice.

[4]We think of the education sector as part of the luxury sector, whereas health care, public safety and well as grocery stores are part of the basic sector

is no scope for confusion. Technology parameters are denoted with Greek letters. Generically, we use the letter x to denote population measures, with super indices specifying subsets of the population. For example, x^{yb} is the measure of young individuals working in the basic sector.

At t_0, the total mass of individuals is one, $x^{yb}+x^{y\ell}+x^{o}=1$, where $x^{yb}=\sum_{i\in\{s,a,f,e,r\}}x^{ybi}$, $x^{y\ell}=\sum_{i\in\{s,a,f,e,r\}}x^{y\ell i}$ and $x^{o}=\sum_{i\in\{s,a,f,e,r\}}x^{oi}$. In the interests of more compact notation, we will also let $x^{i}=x^{ybi}+x^{y\ell i}+x^{oi}$ for $i\in\{s,a,f,e,r\}$ denote the total number of individuals in health state i. Finally, at any point in time let $x=\sum_{i\in\{s,a,f,e,r\}}x^{i}=x^{yb}+x^{\ell b}+x^{o}$ denote the entire living population.

We now describe the dynamics of individuals across health states. The crucial health transitions, and the ones that can, in our model, be affected by mitigation policies are the movements from the susceptible to the asymptomatic state. These are characterized by equations (1)-(5) below. Equation (1) captures the flow of young susceptible individuals working in the basic sector into the asymptomatic state. The number of such workers who catch the virus is their original mass, x^{ybs}, times the number of virus-transmitting interactions they have (the term in square brackets). The three terms in the bracket capture the three sources of infection: from co-workers, from caring for the sick, and from outside the home. The rates of contagion in these different settings depend on how many contagious people a given susceptible basic worker can expect to meet and also on the extent to which people practice social distancing behaviors that reduce the spread of the virus, as measured by the β coefficients. We allow the extent of social distancing to depend on the setting: social distancing at work is denoted β_w, while social distancing outside of work is denoted β_h. We assume that workers in hospitals always take maximum precautions, and that β_w^e is the associated contagion-mitigation parameter.

$$\dot{x}^{ybs}=-\left[\beta_w\ c_w+\beta_h\ c_h+\beta_e\ x^{e}\right]\ x^{ybs} \tag{1}$$

$$\dot{x}^{y\ell s}=-\left[\beta_w\ c_w\ (1-m(t))+\beta_h\ c_h\right]\ x^{y\ell s} \tag{2}$$

$$\dot{x}^{os}=-\beta_h\beta_h^{o}\ \left[\frac{x^{a}+x^{f}}{x-x^{e}}\right]\ x^{os} \tag{3}$$

where

$$c_w=\frac{x^{yba}+(1-m(t))x^{y\ell a}}{x^{ybs}+x^{ybr}+x^{yba}+(1-m(t))\left[x^{y\ell s}+x^{y\ell r}+x^{y\ell a}\right]} \tag{4}$$

$$c_h=\frac{\left[x^{yba}+x^{y\ell a}+x^{ybf}+x^{y\ell f}\right]+\beta_h^{o}\left[x^{oa}+x^{of}\right]}{x-x^{s}}. \tag{5}$$

The first outflow rate in equation (1) given by $\beta_w c_w$ is due to work. It is determined by the

fraction of coworkers who are contagious, c_w, defined in equation (4). We assume that people with symptoms always stay at home (a minimal precaution), and that basic and luxury workers mingle together at work. Thus the fraction of contagious co-workers c_w is the number of asymptomatic workers, $x^{yba} + (1 - m)x^{y\ell a}$ divided by the total number of workers. The social distancing parameter β_w defines the probability a susceptible worker will contract the virus if all his co-workers were asymptomatic carriers of the disease.

The rate at which a young basic worker contracts the virus at home, $\beta_h c_h$, depends on the share of contagious workers in the household, c_h defined in equation (5). Note that both asymptomatic and flu-suffering workers reside at home. We allow for a lower contact rate between young and old within the household via the parameter β_h^o. To the extent that people are especially prone to reduce contact with the old, we will assume $\beta_h^o < 1$. Finally, we assume that caring for those requiring emergency care is a task that falls entirely on basic workers. The risk of contracting the virus from this activity is proportional to the number of hospitalized people, x^e, with proportionality factor β_e measuring the strength of precautions in hospitals.

In parallel to Equation (1), Equation (2) describes infections for the susceptible population working in the luxury sector. It differs from the basic sector only in terms of infections at work. Individuals in this sector work reduced hours when $m > 0$ and thus have fewer work interactions in which they could get infected. Furthermore, workers in the luxury sector do not take care of sick patients in hospitals, and thus the last term in Equation (1) is absent in Equation (2). Finally, Equation (3) displays infections among the old. They only receive infections from interactions at home, and only from the share of the population that carries the disease but is not in emergency rooms. Their rate of infection is further reduced by the factor $\beta_h^o < 1$ capturing the extra precautions society takes when dealing with old, especially vulnerable individuals during the pandemic.

The remainder of the epidemiological block is fairly mechanical and simply describes the transition of individuals though the health states (asymptomatic, flu-suffering, hospitalized, and recovered) once they have been affected. The parameters of these dynamic laws in Equation (6) to Equation (17) are allowed to vary by age. Equations (6) to (8) describes the change in the measure of asymptomatic individuals. There is entry into that state from the newly infected flowing in from the susceptible state (as described above). Exit from this state to those suffering flu-like symptoms occurs at rate σ^{iaf} and to recovered state at rate σ^{iar} for the young and the old, $i \in \{y, o\}$, respectively. Note that someone who recovers at this stage will never know she

contracted the virus or not.

For individuals suffering from the flu, Equations (9) to (11) show that there is entry from the asymptomatic state and exit to the hospitalized state at rate σ^{ife} and to the recovered state at rate σ^{ifr}. Equations (12) to (14) describe the movements of those in emergency care, showing entry from those with flu-like symptoms, and there is exit to death at rate $\sigma^{ied}+\varphi$ and to recovery at rate $\sigma^{ier}-\varphi$, where φ, described below, is a term related to hospital overuse. Equations (15) to (17) displays the evolution of the measure of the recovered population, which features only entry and is an absorbing state. So is death, with the evolution of the deceased population being determined by $\dot{x}^{ybd} = (\sigma^{yed}+\varphi)x^{ybe}$, $\dot{x}^{y\ell d} = (\sigma^{yed}+\varphi)x^{y\ell e}$ and $\dot{x}^{od} = (\sigma^{oed}+\varphi)x^{oe}$. We record them separately from the recovered (who work) since the play no further role in the model.

Finally, Equation (18) shows the extent of overuse of the health care system that has capacity Θ, which we treat as fixed in the time horizon analyzed in this paper. The probability of death conditional on being sick depends on the extent of hospital overuse. In particular, φ measures the amount by which the death rate of the sick rises (and the recovery rate falls) once hospital capacity Θ is exceeded.

$$\dot{x}^{yba} = -\dot{x}^{ybs} - \left(\sigma^{yaf}+\sigma^{yar}\right) x^{yba} \tag{6}$$

$$\dot{x}^{y\ell a} = -\dot{x}^{y\ell s} - \left(\sigma^{yaf}+\sigma^{yar}\right) x^{y\ell a} \tag{7}$$

$$\dot{x}^{oa} = -\dot{x}^{os} - \left(\sigma^{oaf}+\sigma^{oar}\right) x^{oa} \tag{8}$$

$$\dot{x}^{ybf} = \sigma^{yaf} x^{yba} - \left(\sigma^{yfe}+\sigma^{yfr}\right) x^{ybf} \tag{9}$$

$$\dot{x}^{y\ell f} = \sigma^{yaf} x^{y\ell a} - \left(\sigma^{yfe}+\sigma^{yfr}\right) x^{y\ell f} \tag{10}$$

$$\dot{x}^{of} = \sigma^{oaf} x^{oa} - \left(\sigma^{ofe}+\sigma^{ofr}\right) x^{of} \tag{11}$$

$$\dot{x}^{ybe} = \sigma^{yfe} x^{ybf} - \left(\sigma^{yed}+\sigma^{yer}\right) x^{ybe} \tag{12}$$

$$\dot{x}^{y\ell e} = \sigma^{yfe} x^{y\ell f} - \left(\sigma^{yed}+\sigma^{yer}\right) x^{y\ell e} \tag{13}$$

$$\dot{x}^{oe} = \sigma^{ofe} x^{of} - \left(\sigma^{oed}+\sigma^{oer}\right) x^{oe} \tag{14}$$

$$\dot{x}^{ybr} = \sigma^{yar} x^{yba} + \sigma^{yfr} x^{ybf} + (\sigma^{yer}-\varphi) x^{ybe} \tag{15}$$

$$\dot{x}^{y\ell r} = \sigma^{yar} x^{y\ell a} + \sigma^{yfr} x^{y\ell f} + (\sigma^{yer} - \varphi) x^{y\ell e} \tag{16}$$

$$\dot{x}^{or} = \sigma^{oar} x^{oa} + \sigma^{ofr} x^{of} + (\sigma^{oer} - \varphi) x^{oe} \tag{17}$$

$$\varphi = \lambda_o \max\{x^e - \Theta, 0\}. \tag{18}$$

2.2 Economic Block

We start describing how production takes place given a mitigation policy (in Section 2.2.1) and then move to describe the preferences of agents. These take into account longevity and the utility both from being alive and from being in a specific health state. (in Section 2.2.2).

2.2.1 Activity: Technology and Mitigation

There are two production sectors that we label basic and luxury, with workers being unable to move between sectors. The basic sector by assumption is exempted from the mitigation policy, whereas more mitigation (a higher m) reduces the amount of output that can be produced in the luxury sector. We assume a production technology that is linear in labor and thus output in the basic sector is given by the number of young workers employed there;

$$y^b = x^{ybs} + x^{yba} + x^{ybr}. \tag{19}$$

Note that this specification assumes that those individuals carrying the virus but being asymptomatic continue to work.[5] Output in the luxury sector in contrast does depend on the mitigation policy and is given by

$$y^\ell = [1 - m(t)] \left(x^{y\ell s} + x^{y\ell a} + x^{y\ell r} \right). \tag{20}$$

We assume that both sectors produce the same good and thus are perfect substitutes. Under this assumption total output of the single consumption good is determined by:

$$y = y^b + y^\ell. \tag{21}$$

Finally, we also assume that a fixed amount of output, $\eta\Theta$, is spent on emergency health care.

[5] One could instead imagine a policy of tracing contacts of infected people, which would allow the planner to keep some portion of exposed workers at home.

2.2.2 Preferences

The old care about consumption and being alive. Therefore, lifetime utility is given by

$$E\left\{\int e^{-\rho t}\ \left[u^{o}(c_t^{o})+\bar{u}+\widehat{u}_t^{j}\right]\ dt\right\} \tag{22}$$

where expectations are taken with respect to the random timing of death, and where $\bar{u}$ measures the period utility from being alive. Similarly, $\widehat{u}_t^j$ is the intrinsic utility of being in state health j; our calibration will indicate that having flu-like symptoms is bad, and having to be treated in the hospital is very bad. The old value their consumption c_t^o according to the period utility function $u^o(c_t^o)$ at discount the future at rate ρ.

Symmetrically, the young also care about their consumption c_t^y, as well as about their health and about being alive, according to the lifetime utility function:

$$E\left\{\int e^{-\rho t}\ \left[u^{y}\left(c_t^{y}\right)+\bar{u}+\widehat{u}_t^{j}\right]\ dt\right\}, \tag{23}$$

Note that workers who experience flu-like symptoms or are in the hospital do not work. Neither does a fraction m of workers in the luxury sector whose workplaces have been shut down by mitigation policy Therefore in equilibrium young workers will experience different consumption depending on whether they work or not. Thus, expected utility of a worker will depend on the sector in which they work in for two reasons. First, sectors differ in their share of economic activity being shut down (and thus, for the individual worker, the probability of being able to work when healthy enough to do so), and second, the sector will affect the distribution of health outcomes.[6]

3 The Public Sector

In this section we first describe the government policy tools in Section 3.1, and then in Section 3.2 we analyze how public transfers are determined statically to yield a utilitarian period social welfare function. We conclude by posing the dynamic Ramsey optimal policy problem

[6]Note that we have not modeled mortality from natural causes. Over the expected length of the COVID-19 pandemic, mortality from natural causes will be small for both age groups. After the pandemic is over, we will model the shorter remaining expected lifetime for the old in a simple reduced-form by assuming that the old have a higher time discount rate than the young in the post-pandemic period.

which maximizes the time integral of discounted instantaneous social welfare by choice of the optimal time path of mitigation $m(t)$.

3.1 Transfers

The public sector is responsible for two choices: mitigation (shutdowns) $m(t)$ and redistribution to individuals that currently do not or cannot work. We assume that the degree of social interaction within the workplace and outside the workplace (what we called the β's) are determined exogenously outside the model. An alternative interpretation is that the least costly measures of social separation are already in place. What the government chooses is the extent to which it imposes a shutdown of economic activity, via $m(t)$, and how much to transfer to those hurt by shutdowns, those that have fallen sick, and those that have retired. In each instant individuals either work (those healthy enough, not subject to mitigation and not old and thus retired) or do not work. The second policy choice beyond mitigation therefore is redistribution between working and non-working individuals. All workers share a common consumption level c^w and all individuals not working share a common consumption level c^n.[7] The second policy choice is how much to transfer, in each instant t, from the working to the nonworking population. Crucially, we assume that these transfers are costly, denoting by $T(c^n)$ the per-capita cost of transferring consumption c^n to those out of work and without current income. We assume that $T(.)$ is increasing and differentiable.

To simplify notation, denote by $(\mu^n(m,x), \mu^w(m,x))$ the mass of non-working and working people, respectively, as a function[8] of the health population distribution x and current mitigation $m = m(t)$. These are defined as

$$\mu^n(m,x) = x^{y\ell f} + x^{y\ell e} + x^{ybf} + x^{ybe} + m\left(x^{y\ell s} + x^{y\ell a} + x^{y\ell r}\right) \tag{24}$$

$$+ \; x^{os} + x^{oa} + x^{of} + x^{oe} + x^{or} \tag{25}$$

$$\mu^w(m,x) = x^{ybs} + x^{yba} + x^{ybr} + [1-m]\left(x^{y\ell s} + x^{y\ell a} + x^{y\ell r}\right) \tag{26}$$

$$\nu^w(m,x) = \frac{\mu^w(m,x)}{\mu^w(m,x) + \mu^n(m,x)} \tag{27}$$

where $\nu^w(m,x)$ is the share of working individuals in the population. The aggregate resource

[7] This is the allocation chosen by a government that values all individuals equally (equal Pareto weights).
[8] We will suppress the dependence on these variables when there is no room for confusion.

constraint can then be written as

$$\mu^w c^w + \mu^n c^n + \mu^n T(c^n) = y - \eta\Theta = \mu^w - \eta\Theta \tag{28}$$

where we have exploited that $y = \mu^w$ since each productive workers produces one unit of output (and outputs in both sectors are perfect substitutes).

Notice that there are no dynamic consequences of the transfer choice c^n. In particular, this choice has no impact on any health transitions. We can therefore solve a static optimal transfer problem at each t (and given a current level of mitigation $m = m(t)$) that delivers a maximum level of social welfare which we denote $W(m, x)$. We turn to derive this expression now.

3.2 The Instantaneous Social Welfare Function

We now derive the instantaneous social welfare function $W(x, m)$, a necessary ingredient into the optimal mitigation problem of the government. The function $W(x, m)$, assuming that all individuals have log-utility and receive the same social welfare weights, is given by

$$W(x, m) = \max_{c^n, c^w} [\mu^w \log(c^w) + \mu^n \log(c^n)] + (\mu^w + \mu^n)\bar{u} + \sum_{i,j} x^{i,j}\widehat{u}^j \tag{29}$$

where the maximization is subject to the aggregate resource constraint (28). Combining the first order conditions with respect to (c^n, c^w) yields

$$\frac{c^w}{c^n} = 1 + T'(c^n). \tag{30}$$

We can use this relation in the resource constraint to obtain

$$\mu^w (1 + T'(c^n))\, c^n + \mu^n c^n + \mu^n T(c^n) = \mu^w - \eta\Theta \tag{31}$$

Defining net per capita income $\tilde{y}$ and average transfer costs $t(c^n)$ as

$$\tilde{y} = \nu - \frac{\eta\Theta}{\mu^w + \mu^n} \tag{32}$$

$$t(c^n) = \frac{T(c^n)}{c^n} \tag{33}$$

we can rewrite the resource constraint in per-capita terms by dividing by $\mu^w + \mu^n$

$$c^n \left[1 + \nu T'(c^n) + (1-\nu)t(c^n)\right] = \tilde{y} \tag{34}$$

Thus the optimal solution to the government transfer problem is given by the solution to the following system:

$$\begin{aligned} c^n \left[1 + \nu T'(c^n) + (1-\nu)t(c^n)\right] &= \tilde{y} && (35) \\ c^w &= c^n(1 + T'(c^n))) && (36) \end{aligned}$$

for an arbitrary differentiable per capita transfer cost function $T(.)$. We can also express period welfare in per capita terms, using the

$$\begin{aligned} W(x,m) &= [\mu^w + \mu^w]\, w(x,m) && (37) \\ w(x,m) &= \log(c^n) + \nu \log(1 + T'(c^n)) + \bar{u} + \sum_{i,j} \frac{x^{i,j}}{\mu^w + \mu^w} \widehat{u}^j && (38) \end{aligned}$$

where the only endogenous input in the period welfare function c^n solves equation (35). In particular, note that $\mu^w + \mu^w$ is independent of mitigation and thus we can discuss the impact of mitigation on current welfare in terms of the per-capita welfare function $w(x,m)$.

The per capita welfare function shows the basic costs from mitigation m. First, it lowers per-capita income, and through it, the level of consumption. This is the $\log(c^n)$ term in $w(x,)$ which is strictly increasing in net income $\tilde{y}$. In the absence of the cost of transfers, this is the only direct effect of current mitigation. Second, the transfer cost to non-working households distorts risk sharing; this is the second term $\nu \log(1 + T'(c^n))$, which is zero if the marginal transfer cost is zero. Note that an increase in mitigation reduces ν and thus the negative impact of mitigation on current welfare is the more severe, the larger is the marginal cost of transfers. This, ceteris paribus, will reduce the incentives of the government to engage in economically costly mitigation.

To see the intuition for our results most clearly, assume that the transfer cost is linear such that $T(c^n) = \tau c^n$. In this case the optimal allocation is given by:

$$c^w = \tilde{y}$$

$$
\begin{aligned}
c^n &= \frac{\tilde{y}}{1+\tau} \\
w(x,m) &= \log(\tilde{y}) - (1-\nu)\log(1+\tau) + \bar{u} + \sum_{i,j} \frac{x^{i,j}}{\mu^w + \mu^w}\widehat{u}^j
\end{aligned}
$$

Thus the negative economic impact of mitigation is given, in this case, by

$$
\frac{\partial w(x,m)}{\partial m} = \frac{\partial \tilde{y}}{\partial m} + (1+\tau)\frac{\partial \nu}{\partial m} < 0, \tag{39}
$$

since both $\frac{\partial \tilde{y}}{\partial m}$ and $\frac{\partial \nu}{\partial m}$ are negative. In addition, we observe that the larger is the marginal cost of transfers τ the more negative is $(1+\tau)\frac{\partial \nu}{\partial m}$. This is how mitigation and redistribution costs interact: the larger is the marginal cost of redistribution, the larger is the economic cost of mitigation $\frac{\partial w(x,m)}{\partial m}$.

In our quantitative exercises we will assume that the transfer cost function per non-worker is given by the quadratic form $T(c^n) = \frac{\tau}{2}\frac{\mu^n}{\mu^w}(c^n)^2 = \frac{\tau}{2}\left(\frac{1-\nu}{\nu}\right)(c^n)^2$ so that total transfer costs are given by $\mu^n T(c^n) = \mu^w \frac{\tau}{2}\left(\frac{\mu^n c^n}{\mu^w}\right)^2$. This functional form is motivated by the idea that each working household has to transfer $\left(\frac{\mu^n c^n}{\mu^w}\right)$ units of consumption to non-working households. Assuming a quadratic cost of extracting resources from workers, the per-worker cost is thus given by $\frac{\tau}{2}\left(\frac{\mu^n c^n}{\mu^w}\right)^2$.[9] Multiplying this by the total number of workers μ^w gives the total transfer cost. For this specification we obtain as optimal allocations to be inserted in the period welfare function above:

$$
\begin{aligned}
c^n &= \frac{\sqrt{1 + 2\tau\frac{1-\nu^2}{\nu}\tilde{y}} - 1}{\tau\frac{1-\nu^2}{\nu}} && (40) \\
c^w &= c^n(1 + T'(c^n))) = c^n\left(1 + \tau\frac{1-\nu}{\nu}c^n\right) && (41)
\end{aligned}
$$

Note that $\left(1 + \tau\frac{1-\nu}{\nu}c^n\right)$ is the effective price the planner has to pay, on the margin to take one more unit of output from workers to give to non-workers. As transfers and thus non-worker consumption c^n rise, this price effectively rises, reflecting a higher marginal cost to additional redistribution. In addition, since higher mitigation m reduces the share of workers ν and increases the share of non-workers $1-\nu$, the effective price of transfers at the margin

[9] The quadratic form is chosen for analytical convenience, but is not central for our qualitative arguments.

increases with mitigation, and the price rises more the higher is τ.

For future reference, we can also construct expected flow utility for each type

$$\begin{aligned} W^{\ell}(x,m) &= \frac{(x^{y\ell n}+x^{y\ell e}+x^{y\ell r})}{x^{\ell}}\left[(1-m)u(c^{w})+mu(c^{n})+\bar{u}\right] \\ &\quad + \frac{(x^{y\ell f}+x^{y\ell e})}{x^{\ell}}\left[u(c^{n})+\bar{u}-\hat{u}\right] \end{aligned}$$

$$W^{b}(x,m) = \frac{(x^{ybn}+x^{ybe}+x^{ybr})}{x^{b}}\left[u(c^{w})+\bar{u}\right] + \frac{(x^{ybf}+x^{ybe})}{x^{b}}\left[u(c^{n})+\bar{u}-\hat{u}\right]$$

$$W^{o}(x,m) = u(c^{n})+\bar{u}-\frac{(x^{yof}+x^{yoe})}{x^{o}}\hat{u}$$

3.3 Optimal Policy

We now assume there is a government/planner (we use these names as synonymous as there is no time consistency problem) that chooses optimal policy over time by choosing the path of mitigation $m(t)$; the optimal choice of redistribution $T(t)$ is already embodied in the period social welfare function $W(x)$. The policy problem the planner solves is then given by

$$\max_{m(t)} \int_{0}^{\infty} e^{-\rho t} W(x)\ dt. \tag{42}$$

subject to the laws of motion of the population Equation (1) to Equation (18).

In a first step we will approximate the optimal time path of mitigation by functions that are part of the following parametric class of generalized logistic functions of time:

$$m(y) = \frac{\alpha_0}{1+\exp(-\alpha_1(t-\alpha_2))} \tag{43}$$

Here the parameter α_0 controls the level of mitigation at $t = 0$. The parameter α_2 governs when mitigation is reduced, and the parameter α_1 commands how swiftly mitigation is reduced. Note that as $t \to \infty$, $m(t) \to 0$. More generally, the complete characterization of the solution derives from a formal maximization process laid out below.

3.3.1 The Unrestricted Optimal Policy Problem

We have derived the period return function $W(\mathbf{x}, m)$. In addition, the evolution of the state (the distribution of the population by health status $\boldsymbol{x} = (x^{i,j})$) evolves according to the vector-valued equation (summarizing Equations (1) to (17) the paper in a compact form):

$$\dot{\boldsymbol{x}} = G(\boldsymbol{x}, m) \tag{44}$$

To solve for the optimal time path of the scalar mitigation variable is then a straightforward optimal control problem with a multi-dimensional state vector and a one-dimensional control variable. Define the current value Hamiltonian as

$$\mathcal{H}(\boldsymbol{x}, m, \boldsymbol{\mu}) = W(\boldsymbol{x}, m) + \boldsymbol{\mu} G(\boldsymbol{x}, m) \tag{45}$$

where $\boldsymbol{\mu}$ is the vector of co-state variables associated with the population state vector $\boldsymbol{x}$. Necessary conditions at an interior solution for mitigation m are the optimality condition for m

$$\begin{aligned} W_m(\boldsymbol{x}, m) &= -\boldsymbol{\mu} \cdot G_m(\boldsymbol{x}, m) && (46)\\ \dot{\boldsymbol{\mu}} &= \rho\boldsymbol{\mu} - [W_{\boldsymbol{x}}(\boldsymbol{x}, m) + \boldsymbol{\mu} \cdot G_{\boldsymbol{x}}(\boldsymbol{x}, m)] && (47)\\ \dot{\boldsymbol{x}} &= G(\boldsymbol{x}, m) && (48) \end{aligned}$$

The key trade offs with mitigation efforts m are encoded in equation (46). A marginal increase in m entails static economic costs of $W_m(\boldsymbol{x}, m)$ stemming from the loss of output and thus consumption of all individuals in the economy, as encoded in $y^n(m)$. The dynamic benefit is a better change in the population health distribution, as encoded in the vector $G_m(\boldsymbol{x}, m)$. Concretely, as is clear from equations $(1-3)$ an increase in m reduces the outflow of individuals from the susceptible to the asymptomatic state. The value (in units of the objective function) are given by the co-state vector $\boldsymbol{\mu}$.

It should be kept in mind that since $(\boldsymbol{x}, \boldsymbol{\mu})$ are vectors, so are the entities $G_m(\boldsymbol{x}, m) = (G_m^{i,j}(\boldsymbol{x}, m))$ and $W_{\boldsymbol{x}}(\boldsymbol{x}, m) = (W_{x^{i,j}}(\boldsymbol{x}, m))$ and $G_{\boldsymbol{x}}(\boldsymbol{x}, m) = (G^k_{x^{i,j}}(\boldsymbol{x}, m))$ so that equation (46) reads explicitly

$$W_m(\boldsymbol{x}, m) = -\sum_{i,j} \mu^{i,j} G_m^{i,j}(\boldsymbol{x}, m) \tag{49}$$

and a specific row of the vector-valued Section 3.3.1 is given by

$$\dot{\mu}^{i,j} = \rho\mu^{i,j} - \left[W_{\boldsymbol{x}^{i,j}}(\boldsymbol{x}, m) + \sum_k \mu^k G^k_{x^{i,j}}(\boldsymbol{x}, m) \right]. \tag{50}$$

4 Calibration

There is a long list of parameters to specify, most of them epidemiological, and we start with them. We set the population share of the old to be 15%, which is the current fraction of the US population aged 65 and above.

There are twelve σ parameters to calibrate, describing transition rates for disease progression, six for each age. These describe the chance of moving to the next worse health status and the chance of recovery at the three infectious stages: asymptomatic, flu-suffering, and hospitalized. We assume that young and old exit each stage at the same rate, but potentially differ in terms of the share of these exits that are into recovery. In particular, the old will be much more likely to require hospital care conditional on developing flu, and slightly more likely to die conditional on being hospitalized.

Putting aside these differences by age for a moment, the six values for σ are identified from the following six target moments: the average duration of time individuals spend in the asymptomatic (contagious but asymptomatic), flu-suffering (relatively mild symptoms), and emergency-care states, and the relative chance of recovery (relative to disease progression) in each of the three states. Following the literature on COVID-19 models we set the three durations to 5.2, 10, and 8 days, with these durations common across age groups. The exit rate to recovery from the asymptomatic state defines the number of asymptomatic cases of COVID-19 and is an important but highly uncertain parameter. We assume that asymptomatic recovery and progression to the flu-suffering state are equally likely.[10]

We let the relative recovery rates from the flu-suffering and emergency care states vary with age, to reflect the fact that infections in older individuals are more much likely to require hospitalization, and hospitalizations are also more likely to lead to death. We set the recovery

[10] Given that the asymptomatic state has roughly half the duration of the flu state, this implies that roughly half of infected agents in the model will be asymptomatic. Recall that in a random sample in Iceland, half of the positive subjects reported no symptoms.

rate from flu-suffering to 96% for the young, and to 75% for the old, based on evidence from Table 1 of the Imperial College study. Similarly, given evidence on differential mortality rates, we set the recovery rates from the emergency care state to 95% for the young and to 80% for the old (assuming no hospital overuse).

Given the σ parameters, the β parameters determine the rate at which contagion grows over time. We set $\beta_e = 0.01$, implying that a very small share of overall transmission occurs in hospitals. The values of β_w and β_h determine the overall basic reproduction number R_0 value for COVID-19 and the share of disease transmission that occurs at work versus in non-work settings. Estimates for R_0 for COVID-19 absent additional social distancing measures are in the range of 2 to 2.5. However, we will focus on starting our simulations assuming fairly severe social distancing measures are already in place, but no economic mitigation ($m = 0$). To start with we will assume this implies $R_0 = 1.40$. Mossong et al. (2008) find that 35 percent of transmission happens in workplaces and schools, with the rest in the home and social settings. We use these targets to pin down choices for β_w and β_h as follows.

The basic reproduction number R_0 is the number of people infected by a single asymptomatic person. For a single young person, assuming everyone else in the economy is susceptible (a close approximation to the initial condition we will use) and zero mitigation ($m = 0$), R_0 is given by

$$R_0^y = \frac{\beta_w + \beta_h}{\sigma^{yar} + \sigma^{yaf}} + \frac{\sigma^{yaf}}{\sigma^{yaf} + \sigma^{yar}} \frac{\beta_h}{\sigma^{yfr} + \sigma^{yfe}} + \frac{\sigma^{yaf}}{\sigma^{yaf} + \sigma^{yar}} \frac{\sigma^{yfs}}{\sigma^{yfe} + \sigma^{yfr}} \frac{\beta_e}{\sigma^{yer} + \sigma^{yed}}$$

The logic is that this individual will spread the virus while asymptomatic, flu-suffering, and sick – the three terms in the expression. They expect to be asymptomatic for $(\sigma^{yar} + \sigma^{yaf})^{-1}$ days, flu-suffering (conditional on reaching that state) for $(\sigma^{yfr} + \sigma^{yfe})^{-1}$ days, and hospitalized (conditional on reaching that state) for $(\sigma^{yer} + \sigma^{yed})^{-1}$ days. The chance they reach the flu-suffering state is $\frac{\sigma^{yaf}}{\sigma^{yaf}+\sigma^{yar}}$ and the chance they reach the emergency room is the product $\frac{\sigma^{yaf}}{\sigma^{yaf}+\sigma^{yar}} \frac{\sigma^{yfs}}{\sigma^{yfe}+\sigma^{yfr}}$. While asymptomatic, they spread the virus at both work and at home, and pass the virus on to $\beta_w + \beta_h$ susceptible individuals per day. While flu-suffering, they stay at home and pass the virus to β_h people per day. While sick they pass it to β_e people per day in hospital.

The reproduction number for an old asymptomatic person is

$$R_0^o = \frac{\beta_h}{\sigma^{oar} + \sigma^{oaf}} + \frac{\sigma^{oaf}}{\sigma^{oaf} + \sigma^{oar}} \frac{\beta_h}{\sigma^{ofr} + \sigma^{ofe}} + \frac{\sigma^{oaf}}{\sigma^{oaf} + \sigma^{oar}} \frac{\sigma^{ofs}}{\sigma^{ofe} + \sigma^{ofr}} \frac{\beta_e}{\sigma^{oer} + \sigma^{oed}}$$

where this formula is similar to the one for the young, except that it recognizes the old pass the virus on less because they do not work. At the same time, however, because the old are less likely to recover once infected, they potentially carry the virus for a longer time, inducing more transmission in hospitals.

For the population as a whole, the overall R_0 is a weighted average of these two group-specific values

$$R_0 = \mu_y R_0^y + (1 - \mu_y) R_0^o$$

where μ^y is the fraction of the population that is young.

The share of total transmission that occurs in the workplace from a randomly drawn newly asymptomatic individual is then given by

$$\frac{\text{workplace transmission}}{\text{all transmission}} = \mu^y \left(\frac{\beta_w}{\sigma^{yar} + \sigma^{yaf}} \right) \frac{1}{R_0}$$

Given these two equations, we set β_w and β_h to hit an $R_0 = 2.25$ and a share of workplace transmission equal to 35 percent. For most of our simulations we then scale down both β's symmetrically to reduce the initial R_0 to 1.40.

To set the value of life $\bar{u}$ we follow the value of a statistical life VSL approach. The Environmental Protection Agency and the Department of Transportation assume a value of \$11.5 million (see Greenstone and Nigam 2020). This is a relatively high value, relative to VSL values used in other contexts. Assuming an average of 37 residual life years discounted at a 3 percent rate, this translates to an annual flow value of $\$515,000$, which is 11.4 times yearly per capita consumption in the United States.

To translate this into a value for $\bar{u}$ we use the standard value of a statistic life calculation,

$$VSL = \frac{dc}{dr}|_{E[u]=k} = \frac{\ln(\bar{c}) + \bar{u}}{\frac{1-r}{\bar{c}}}$$

where $\bar{c}$ is average per capita model consumption, and r is the risk of death. Setting $VSL = 11.4\bar{c}$ and $r = 0$ gives $\bar{u} = 11.4 - \ln\bar{c}$. Note that this implies an easily interpretable trade-off between mortality risk and consumption. For example, we can ask what reduction in consumption leads an individual indifferent to facing a 1 percent risk of death. The answer is the solution m to

$$\ln(\bar{c}\,(1-m)) + 11.4 - \ln\bar{c} = 0.99\,(\ln(\bar{c}) + 11.4 - \ln\bar{c})$$

which is $m = 1 - \exp(-0.01 \times 11.4) = 10.8\%$.

As another way to get a feeling for what our choice for the value of a statistical life implies, suppose we were to contemplate a shut down that would reduce consumption for six months by 25 percent. By how much would this shut down have to reduce mortality risk for an agent with 10 expected years of life to prefer the shutdown to no shutdown? The answer is the solution x to

$$\frac{1}{20}\ln(1-0.25) + \frac{19}{20}\ln(1) + 11.4 = (1-x)11.4$$

which is 0.13 percent.

For the disutility of having flu, we define $\hat{u}^f$ as

$$\hat{u}^f = -0.3\,(\ln(\bar{c}) + \bar{u})$$

following Hong et al. (2018). We set $\hat{u}^e = -(\ln(\bar{c}) + \bar{u})$, so that the flow value of being in hospital is equal to the flow value of being dead (zero).

We set emergency health care capacity parameter Θ to 0.025 percent, reflecting around $80,000$ ICU beds. Because the cost of a day in intensive care is around \$7,500, we set $\eta = 50$, so emergency care consumes about 1.5 percent of pre-COVID output.[11] We set the parameter λ_0 such that, absent economic mitigation, the recovery rate of the old in emergency care at the peak of the epidemic is half of its value when capacity is not exceeded.

The remaining parameters have to do with the economic side of the model. We set the discount rate ρ equal to 3 percent per year. It is hard to gauge the size of the basic sector. At a minimum, it seems hard to do with food production and distribution, health care, utilities,

[11] Total healthcare spending in the United States is 18 percent of GDP. Of this, around 1/3 is spending on hospitals.

police and fire, basic financial services and liquor stores. For now we assume that the basic sector is 30 percent of the economy.

We adopt the quadratic formulation of transfer costs. We pick a value for τ using estimates for the excess burden of taxation, which suggest that raising an extra dollar in revenue at the margin (which can be used to increase consumption for non-workers) has a cost on taxpayers of around \$1.38 (Saez, Slemrod and Giertz, 2012). This suggests $\tau \frac{1-\nu}{\nu} c^n = 0.38$. Given the first order condition above, this suggests an optimal redistribution scheme would imply $c_n/c_w = 1/1.38 = 0.72$ in pre-COVID times. Moreover, given $\eta\Theta = 0.0125$, $\tau \frac{1-\nu}{\nu} c^n = 0.38$ and $\nu = 0.85$, section 3.2 implies $\tau = 3.47$.

For the time path of mitigation, our baseline simulation, designed to approximate current US policy, will assume $m = 0.5$ for 100 days followed by $m = 0$ thereafter. This path is implemented in the context of the mitigation function (eq. 43) by setting $\alpha_0 = 0.5$, $\alpha_1 = -0.3$, and $\alpha_2 = 100$. As an initial condition, we start the model off with 0.5 percent of the population in the asymptomatic state, 0.5 percent in the flu state, and the remaining 99 percent susceptible.

We simulate the model for 500 days, which in all the cases we have explored is sufficient for the pandemic to have run its course. When evaluating welfare, we discount utility for old at young at a three percent annual rate for the first 500 days. After than we apply different discount factors to the two groups to compute remaining lifetime utility in the final steady state, as a simple way to take into account shorter remaining life expectancy for the old. In particular, we think of the typical young person being 32.5 years old with 47.5 expected years to live, and the typical old person being 72.5 with 14 years to live, where these life expectancies are taken from actuarial life tables. Given these values, and a pure discount rate of three percent, adjusted discount rates that incorporate differential expected longevity are 4 percent for the young, and 10 percent for the old.

Table 1: Epidemiological Parameter Values

Behavior-Contagion (with social distancing)

β_w	infection at work	35% of infections at work	0.11
β_h	infection at home	$R_0 = 1.40$	0.09
β_e	infection in hospitals		0.01
β_h^o	lower interaction for old		1.0

Disease Evolution

σ^{yaf}	rate for young asymptomatic into flu	50% flu, 5.2 days	$\frac{0.5}{5.2}$
σ^{yar}	rate for young asymptomatic into recovered		$\frac{0.5}{5.2}$
σ^{oaf}	rate for old asymptomatic into flu	50% flu, 5.2 days	$\frac{0.5}{5.2}$
σ^{oar}	rate for old asymptomatic into recovered		$\frac{0.5}{5.2}$
σ^{yfe}	rate for young flu into emergency	4% hospitalization, 10 days	$\frac{0.04}{10}$
σ^{yfr}	rate for young flu into recovered		$\frac{0.96}{10}$
σ^{ofe}	rate for old flu into emergency	25% hospitalization, 10 days	$\frac{0.25}{10}$
σ^{ofr}	rate for old flu into recovered		$\frac{0.75}{10}$
σ^{yed}	rate for young emergency into dead	0.15% mortality, 8 days	$\frac{0.05}{8}$
σ^{yer}	rate for young emergency into recovered		$\frac{0.95}{8}$
σ^{oed}	rate for old emergency into dead	5.0% mortality, 8 days	$\frac{0.20}{8}$
σ^{oer}	rate for old emergency into recovered		$\frac{0.80}{8}$

Table 2: Economic Parameters

ρ	pure discount rate	3.0% per year	$\frac{0.03}{365}$
ρ_+^y	effective discount rate of young	4.0% per year	$\frac{0.01}{365}$
ρ_+^o	effective discount rate of old	10% per year	$\frac{0.07}{365}$
$\bar{u}$	value of life	11.4× consumption p.c.	
$\widehat{u}^f$	disutility of flu	lose 30% of baseline utility	
$\widehat{u}^e$	disutility of emergency care	lose 100% of baseline utility	
$\frac{x^{yb}}{x^{yb}+x^{y\ell}}$	size of basic sector	30%	0.30
$\frac{x^o}{x^y+x^o}$	share of old	15%	0.15
τ	transfer cost	\$0.38 burden of excess taxation	3.47
α_0	initial share mitigated	50%	0.5
α_1	speed of mitigation		−0.3
α_2	time mitigation begins	100 days	100
Θ	hospital capacity	80,000 beds	0.00025
λ_o	impact of overuse on mortality	double mortality at peak	

5 Findings

We start by describing model outcomes under what we think of as the policies currently in place in the United States. We then turn to optimal mitigation in the next section.

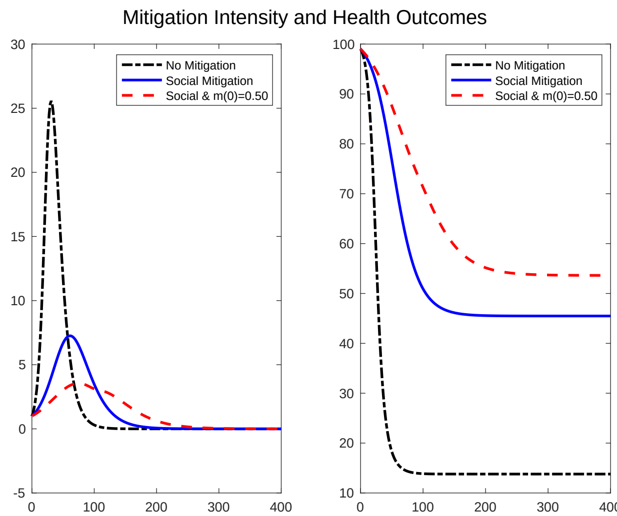

Figure 1: Left panel: share of the population infected (asymptomatic + flu + hospitalized). Right panel: share of population never infected (susceptible).

The left-hand panel of Figure 1 shows dynamics for the fraction of the population infected (asymptomatic plus flu plus hospitalized) in three scenarios: (1) β_h and β_w at pre-social distancing levels, and no economic mitigation ($m = 0$), implying $R_0 = 2.25$; (2) β_h and β_w reduced to our baseline values, with no economic mitigation, implying $R_0 = 1.40$; and (3) β_h and β_w reduced to our baseline values, with 50 percent economic mitigation ($m = 0.5$) for 100 days, implying $R_0 = 1.23$. We think of this third scenario as roughly approximating the current policy in place.

Absent any mitigation, the virus spreads rapidly, and at its peak, less than 50 days from date zero, over a quarter of individuals are infected. Social distancing significantly dampens and postpones the peak in infections. Imposing moderate economic mitigation on top of these

baseline social distancing measures further flattens the curve and pushes the peak out to around 90 days, at the cost of higher infection rates at later dates.

The right panel of the plot shows the share never infected (i.e., susceptible) under the same three scenarios. The key message is that more aggressive mitigation measures do not just flatten the curve: they also reduce the total number of infections. The logic is that in the SIR class of models, the growth rate of infections depends not just on how many people are infected, but also on the relative shares of susceptible versus recovered individuals in the non-infected population. More aggressive mitigation measures slow the spread of infection, such that infections peak later. But delaying the peak in infections gives time for more people to recover and develop immunity, which slows infection growth. The result is that we converge to a steady state in which a larger share of individuals are never infected.

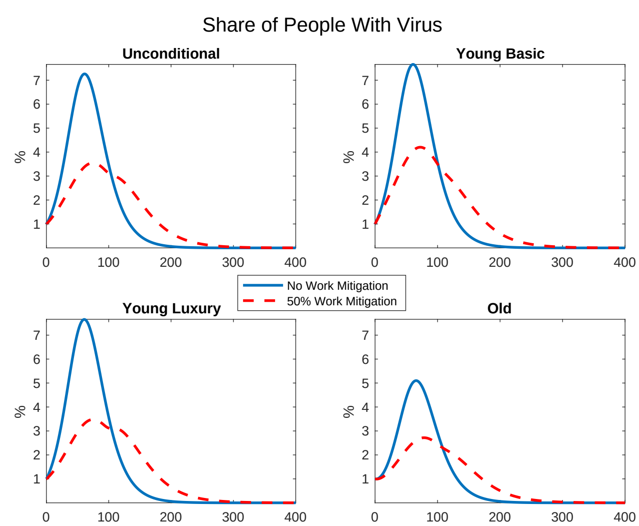

Figure 2: Share of each group infected (asymptomatic + flu + hospitalized).

The next figure (Figure 2) breaks out the share infected by type. Here we report only the economy with social distancing already in place, and the economy with economic mitigation on top of social distancing. Thus, the only difference between the two lines plotted is that $m = 0$ in the blue economy, and $m = 0.5$ in the red economy. Note that, absent economic mitigation,

basic and luxury sector workers are infected at nearly identical rates, while the old – who do not face exposure at work – experience a lower rate of infection. Economic mitigation reduces infection rates for all three types. For the two types of workers, the effect is slightly larger for luxury workers – the type that stays home. But all three groups benefit to a surprisingly similar extent, reflecting the fact that lower virus spread at work means fewer infected people outside work, and thus less new infection at home.

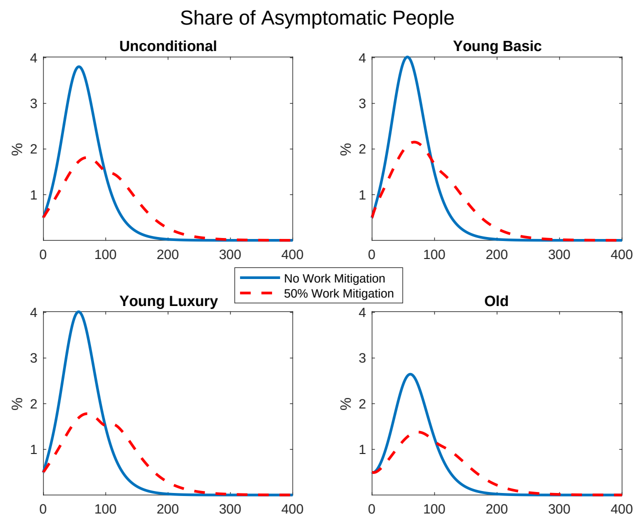

Figure 3: Share of each group infected but asymptomatic

The next three plots break down the infected populations into those who are asymptomatic, those with flu symptoms, and those who are hospitalized. The key thing to note here is that while a smaller share of the old develop mild symptoms – reflecting a lower infection rate – a much larger share end up hospitalized (Figure 4). This is because the old are much more likely to experience progression from flu to more serious symptoms. The red horizontal line plots hospital capacity, Θ. Clearly capacity is drastically exceeded for many days under both scenarios.

Figure 6 shows the evolution of share of the initial population that has died from COVID-19. The virus kills around 0.4 percent of the population with social distancing in place. Imposing

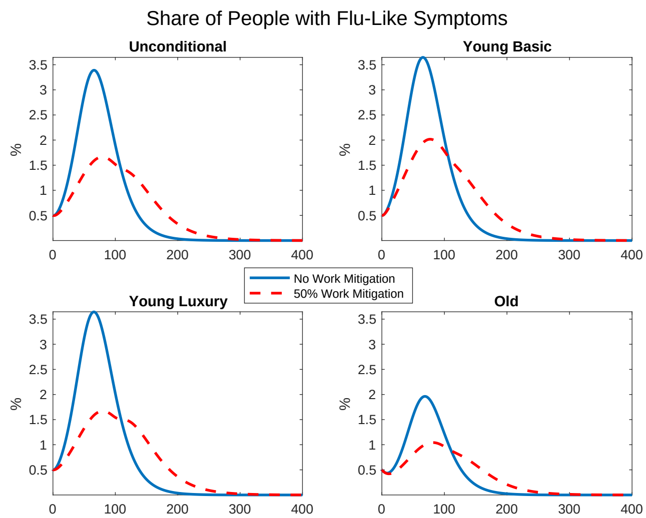

Figure 4: Share of each group with flu symptoms

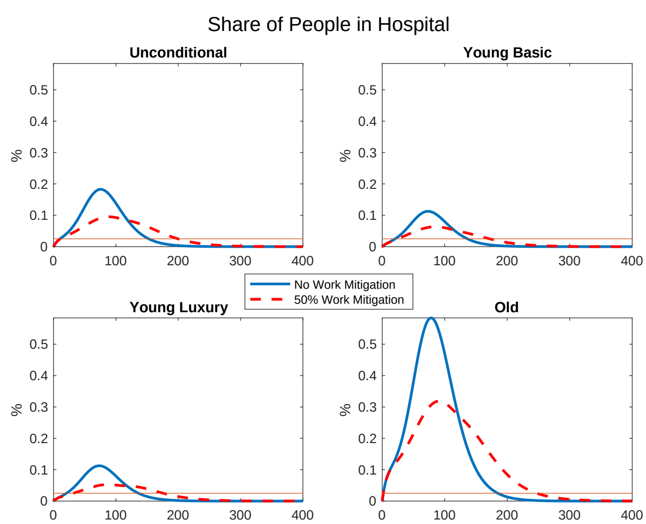

Figure 5: Share of each group hospitalized.

economic mitigation on top of that reduces mortality to around 0.25 percent of the population, which amounts to 817,000 people. Note that the virus kills a much larger share of the old, reflecting a greater likelihood that older infected individuals require hospitalization, and a smaller chance of recovery post hospitalization.

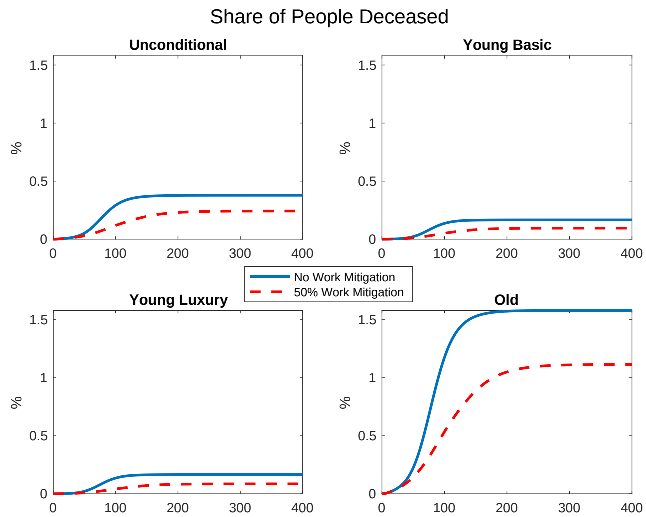

Figure 6: Share of each group deceased.

Figure 7 plots the dynamics of consumption for workers and non-workers through the course of the pandemic. Recall that in this economy all workers enjoy the same consumption level, independent of sector, and the government provides equal consumption via transfers to all non-workers, irrespective of whether they are not working because they are old, sick, or asked to stay home because of economic mitigation. The four panels correspond to four different economies. In the top two panels, we assume use our baseline value for τ, which implies that it is costly for the planner to redistribute from workers to non-workers. In the bottom two panels, we set $\tau = 0$, so that the planner can redistribute freely. In that case, the planner equates consumption between workers and non-workers at each date.

The left two panels describe the evolution of consumption absent economic mitigation ($m = 0$). There is a mild consumption recession in this case, reflecting the fact that infected people with symptoms are assumed to stay home rather than go to work. The right two panels describe the cost in terms of lost consumption from economic mitigation. Given that we are

shutting down 50 percent of a sector that ordinarily accounts for 70 percent of economic activity, we see consumption levels that are around two thirds of those in the unmitigated economy. Note also that the cost of economic mitigation is born disproportionately by non-workers: the ratio of non-worker to worker consumption declines (from two thirds to one half) during the mitigation phase. This reflects our assumption that extracting resources to redistribute from workers becomes ever harder the more the planner wants to tax each worker. To avoid very large redistribution costs, the planner optimally chooses to reduce insurance during the mitigation phase.

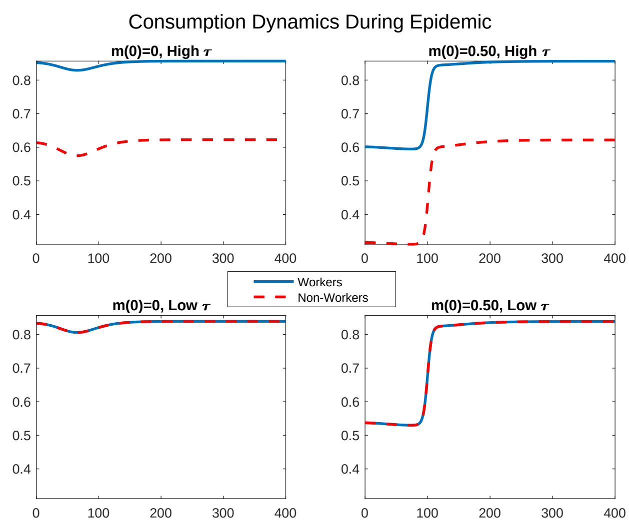

Figure 7: Consumption paths. Top two panels, $\tau = 0.347$. Bottom two panels, $\tau = 0.001$. Left two panels, $m = 0$. Right two panels, $m = 0.5$ for 100 days, then $m = 0$.

Next we report the expected welfare gains and losses for each type of individual for various assumptions about the level of economic mitigation and the parameter τ that indexes the cost of redistribution. In particular, we consider three mitigation levels: $m = 0.5$ (our baseline used to construct the previous plots), $m = 0.75$, and $m = 0.25$. In each case we assume mitigation is in place for 100 days. The welfare calculation asks: what percent of consumption would a person pay every day for the rest of his life to move from the economy with $m = 0$ to $m = 0.5$ (or $m = 0.75$ or $m = 0.25$) for 75 days. We report results for our baseline value for τ (3.47)

Table 3: Welfare Gains (+) or Losses (-) From Mitigation

Mitigated Share	75%		50%		25%	
Transfer Cost (τ)	3.47	0.001	3.47	0.001	3.47	0.001
Young Basic	0.25%	0.08%	0.39%	0.29%	0.29%	0.24%
Young Luxury	-0.25%	0.17%	0.17%	0.40%	0.24%	0.33%
Old	3.48%	4.39%	3.59%	4.33%	2.32%	2.75%

and for a case in which redistribution is costless ($\tau = 0$).

The first clear message from Table 3 is that economic mitigation is generally welfare improving. In particular all types experience expected welfare gains from either 25 percent or 50 percent mitigation policies.

The second message is that the welfare gains are much larger for the old than for the young. For example, in our baseline case ($m = 0.5$ and $\tau = 3.47$) the old gain over 3.5 percent of consumption, while the young gain less than 0.5 percent. The reason the gains are much larger for the old is simply that the old face a much higher likelihood of being killed by the virus.

The third key message is that the cost of redistribution matters. In particular, when redistribution is costless, young luxury workers gain more from mitigation than young basic workers. The logic is that both types of workers suffer identical consumption losses, but the luxury workers benefit more from reduced infection at work. However, when redistribution is costly, the situation is reversed. Now young luxury workers benefit less than young basic workers, because they suffer larger consumption losses from economic mitigation. The reason is that when mitigation is increased, the planner needs to redistribute from a smaller pool of workers toward a larger pool of non-workers. Given convex costs of extracting additional resources from workers, this induces the planner to reduce insurance, translating into a larger consumption gap between workers and non-workers.

We now briefly discuss a few factors that shape these welfare calculations. First, the overall level of the welfare numbers is sensitive to several choices. A key one is the value of a statistical life: a higher value would make mitigation even more attractive. Second, if we assumed lower recovery rates at different stages of infection, or a higher mortality rate at the hospital stage,

agents would perceive a greater risk of death, and again be more willing to sacrifice consumption to avoid that risk. Third, in our model, when a shutdown raises non-employment and reduces consumption, there is no upside in households' utility functions from more leisure. In the analysis of optimal shutdowns in Eichenbaum et al. (2020), the fact that households experience reduced disutility from labor supply when economic activity is taxed makes the utility cost of reduced consumption much smaller.

5.1 Optimal Policy

The mitigation policies we have chosen to date were not chosen optimally. We now turn to explore the optimal time path for economic mitigation. To start, we optimize over two parameters in our parametric process for m_t: α_0, which controls the initial level of mitigation, and α_2, which controls when mitigation ends.

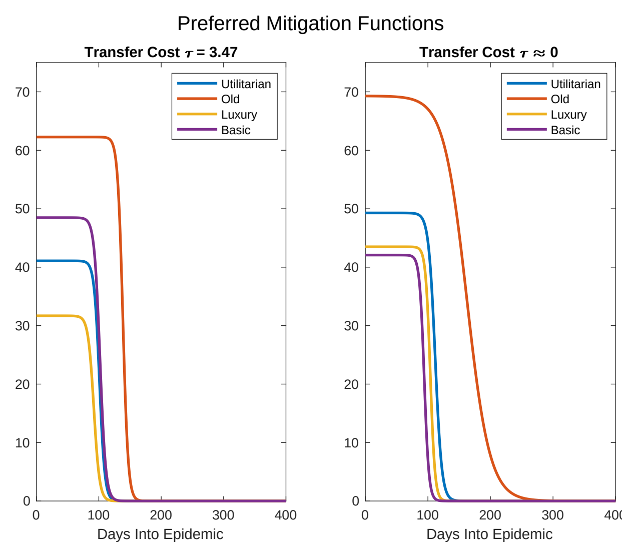

Figure 8: Preferred intensity and duration of lock downs Left panel costly redistribution. Right panel costless redistribution

Figure 8 describes the preferred policies within this class. The left-panel describes optimal policies under our baseline policy, with $\tau = 3.47$. The blue line is the policy chosen by a

Table 4: Welfare Gains (+) or Losses (-) From Preferred Mitigation, $\tau = 3.47$

	Utilitarian	Old	Young Luxury	Young Basic
Young Basic	0.38%	0.27%	0.33%	0.39%
Young Luxury	0.22%	-0.23%	0.25%	0.17%
Old	3.20%	3.91%	2.59%	3.42%

utilitarian planner, who weights expected utility of each type in proportion to their date 0 population shares. The other lines describe the policies preferred by each of the three different types. The right-panel corresponds to a case in which there are no costs to mitigation.

There are clearly large different across types in terms of what fraction of the economy they would like to see shutdown, and for how long. These differences are especially pronounced in the baseline model with costly redistribution. The old (15% of the population) would like to see over 60% of the luxury shut down, and for it to remain shut down for amount 150 days. In contrast, young luxury workers (60% of the population) would prefer to see only half as much of the luxury sector shuttered, and for less than 100 days. Basic sector workers prefer an intermediate policy. Balancing these very different preferences, a utilitarian planner choose to shut down a little over 40% of the luxury sector (28% of the whole economy) for around 100 days.

When redistribution is costless, policy preferences are quite different. Now the old prefer even more mitigation, because they do not have to worry about reduced insurance during a shutdown. Young luxury and young basic sector workers now have very similar policy preferences. Interesting, basic sector workers are the only group that favors less mitigation when there is better insurance. Intuitively when redistribution is costless, more mitigation translates into more expected redistribution from this group toward everyone else.

The next two tables (Tables 4 and 5) describe expected welfare gains, relative to a no economic mitigation baseline, under each of the policies described in Figure 8. The columns of each table identify the policy in place. The rows report expected welfare for each type.

The first takeaway from these tables is that the old experience large welfare gains from any of these policies, irrespective of the cost of redistribution, while welfare gains or losses for the young are much smaller. Second, welfare gains for young luxury workers are always

Table 5: Welfare Gains (+) or Losses (-) From Preferred Mitigation, $\tau \approx 0$

	Utilitarian	Old	Young Luxury	Young Basic
Young Basic	0.27%	-0.11%	0.29%	0.30%
Young Luxury	0.39%	0.04%	0.40%	0.40%
Old	4.48%	5.30%	4.08%	3.78%

smaller than for young basic workers when redistribution is costly, while the pattern is reversed when redistribution is costless. Third, when redistribution is costly, the policy that is welfare maximizing for the old is actually welfare-reducing (relative to no mitigation) for young luxury workers.

6 Conclusion

In this paper we have extended a standard epidemiological model of disease progression to include heterogeneity by age, and multiple sources of disease transmission, and incorporated it into a multi-sector economic model in which workers differ by sector (basic and luxury) as well as by health status. We have studied optimal epidemic mitigation policies and have argued that costly redistribution reduces the desire of the government to engage in mitigation policies. Our results also starkly illustrate how unevenly the welfare gains and losses from economic mitigation are distributed across different segments of society.

References

ATKESON, A. (2020): "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios," Working Paper 26867, National Bureau of Economic Research.

BAYER, C. AND M. KUHN (2020): "Intergenerational ties and case fatality rates: A cross-country analysis," London, Centre for Economic Policy Research DP14519.

BERGER, D., K. HERKENHOFF, AND S. MONGEY (2020): "An SEIR Infectious Disease Model with Testing and Conditional Quarantine," BFI Working Paper NO. 2020-25.

EICHENBAUM, M., S. REBELO, AND M. TRABANDT (2020): "The Macroeconomics of Epidemics'. London, Centre for Economic Policy Research," CEPR DP 14520.

FANG, H., L. WANG, AND Y. YANG (2020): "Human Mobility Restrictions and the Spread of the Novel Coronavirus (2019-nCoV) in China," Working Paper 26906, National Bureau of Economic Research.

FLAXMAN, S., S. MISHRA, A. GANDY, ET AL. (2020): "Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries," Imperial College London (2020) doi.

GREENSTONE, M. AND V. NIGAM (2020): "Does Social Distancing Matter?" (March 30, 2020). University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2020-26. Available at SSRN: https://ssrn.com/abstract=3561244 or http://dx.doi.org/10.2139/ssrn.3561244.

HALL, R. E., C. I. JONES, AND P. J. KLENOW (2020): "Trading Off Consumption and COVID-19 Deaths," Stanford University Working Paper.

HONG, J. H., J. PIJOAN-MAS, AND J.-V. RIOS-RULL (2018): "Health and Inequality," Work in Progress. First Version, 2016.

KRUEGER, D., H. UHLIG, AND T. XIE (2020): "Macroeconomic Dynamics and Reallocation in an Epidemic," Draft, University of Chicago.

MOLL, B., G. KAPLAN, AND G. VIOLANTE (2020): "Pandemics According to HANK," Virtual Presentation on March 31, 2020.

MOSSONG, J., N. HENS, M. JIT, P. BEUTELS, K. AURANEN, R. MIKOLAJCZYK, ET AL. (2008): "Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases," *PLoS Med*, 5, https://doi.org/10.1371/journal.pmed.0050074.

TOXVAERD, F. (2020): "Equilibrium Social Distancing," Tech. rep., Cambridge-INET Working Paper Series No: 2020/08 Cambridge Working Papers in Economics: 2021.