PIER Working Paper

20-014

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

ANDREW GLOVER
Federal Reserve Bank of Kansas City

JONATHAN HEATHCOTE
Federal Reserve Bank of Minneapolis and CEPR

DIRK KRUEGER
University of Pennsylvania NBER & CEPR

JOSÉ-VÍCTOR RÍOS-RULL
University of Pennsylvania, UCL, CAERP, CEPR and NBER

Orginal March 31, 2020
Revised July 21, 2020

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

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

July 21, 2020

First version, March 31 2020

Abstract

To slow COVID-19, many countries have shut down part of the economy. Older individuals have the most to gain from slowing virus diffusion. Younger workers in sectors that are shuttered have most to lose. In this paper, we build a model in which economic activity and disease progression are jointly determined. Individuals differ by age (young, retired), by sector (basic, luxury), and health status. Disease transmission occurs in the workplace, through consumption, at home, and in hospitals. We study the optimal economic mitigation policy for a government that can redistribute across individuals, but where redistribution is costly. Optimal redistribution and mitigation policies interact, and more modest shutdowns are optimal when redistribution is more costly. We find that the shutdowns that were implemented in mid-April were too extensive, but a partial shutdown should remain in place through the fall. A deeper and longer shutdown is preferred if a vaccine is imminent.

Keywords: COVID-19; Economic Policy; Redistribution

1 Introduction

Health pandemics such as the COVID-19 crisis unfold rapidly, and individual behavior has large externalities on others in society. The central debate about the appropriate economic policy response to a pandemic is how aggressively to restrict economic activity in order to slow 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 response is that the benefits and costs of "lock-downs" are large and very unequally distributed. Thus different groups prefer very different policies.

Standard epidemiological models describing the dynamic of a pandemic assume a representative agent structure, in which all households face the same trade-off between restrictions on social interaction that slow virus transmission but which also depress economic activity. In reality, for COVID-19 the benefits of slower viral transmission accrue disproportionately to older households, who face much higher risk of serious illness or death from infection. In contrast, the costs of reduced economic activity are disproportionately borne by younger households bearing the brunt of lower employment. For these younger households, the costs of mitigation policies depend crucially on their sector of work. Sensible lock-down policies designed to reduce viral spread will naturally focus on reducing activity in sectors in which there is a social aspect to consumption and sectors that produce goods or services perceived to be non-essential. For example, restaurants and bars are likely to be closed first. Because workers cannot easily reallocate across sectors, this implies that lock-down policies will involve extensive redistribution among 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.

One way to try to build a coalition in favor of mitigation efforts is to use redistributive public policies to reduce the costs to those whose jobs are threatened by shutdowns. However, redistribution is costly in practice. The more costly is redistribution, the larger and more unequal will be the welfare costs of mitigation measures. Understanding and quantifying these distributional trade-offs is central to thinking about how to respond to pandemics.

In this paper, we build and then quantitatively implement a model that implements 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 where consumption, production, caring for the sick and purely social interactions determine health transitions during the epidemic, and (iii) a government with tools for mitigation and redistribution, as well as a desire for social insurance, and (iv) costs to redistribution which curbs the incentive of the government to equalize the burden of the pandemic across individuals.

We distinguish between three types of people: young workers in a basic sector, young workers in a luxury sector, and old retired people. The output of workers in the two sectors is combined to produce a single final consumption good. Workers are immobile across sectors. The output of the basic sector is assumed to be so essential that it will not make sense to reduce employment and output in that sector in order to reduce the spread of the disease. In contrast, the policy maker has a potential incentive to shut down part of the economic activity in the luxury sector in order to reduce the rate at which infection spreads.

The epidemiological model builds on a standard Susceptible-Infectious-Recovered (SIR) diffusion framework. We label our variant a *SAFER* model, reflecting the progression of individuals through a sequence of possible health states. Individuals start out as susceptible, S (i.e., healthy, but vulnerable to infection), and can then become infected but asymptomatic, A; infected with fever-like symptoms, F; infected and needing emergency hospital care, E, recovered, R (healthy and immune), or dead. The transition rates between these states vary with age; the old are much more likely to experience adverse health outcomes conditional on being infected.

At the heart of the model is the two-way interaction between the distribution of health and economic activity. We model virus transmission from co-workers in the workplace, from co-consumers in the marketplace, from friends and family at home, and from the sick in hospitals. Because they do not work, the old do not face direct exposure at work, but virus transmission in the workplace indirectly increases infection rates in other settings. Our three different infected subgroups spread the virus in very different ways: the asymptomatic are unlikely to realize they are contagious and will continue to work and to consume; those with a fever will stay at home and only infect family members, while those in hospital care may pass the virus to health care workers.

The government uses a utilitarian social welfare function and has at its disposal two policy levers 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 asymptotically infected, but it reduces to zero the market income of some workers in the luxury sector. Second, the planner chooses how much income to redistribute from those working to those that are not, because they are old, because they are unwell, or because their workplaces have been closed due to mitigation. Redistribution is desirable because of the utilitarian social welfare function. Crucially, we 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 a static social planner problem. More mitigation is associated with lower aggregate consumption and more consumption inequality across workers, the more so the costlier is redistribution. This in turn affects adversely the dynamic incentives for mitigation, implying that a government facing more costly redistribution will dynamically choose less mitigation.

In the context of the model with these trade-offs, we then compute optimal paths for mitigation. We find that a planner who prioritizes the old chooses extensive and prolonged mitigation, as the old are highly vulnerable to contracting and dying from the disease. A planner who prioritizes workers in the luxury sector subject to shut-downs chooses a much milder and shorter mitigation path, as the economic costs of forgone income and thus consumption dominate the health benefits 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 this cost is, the more moderate is optimal mitigation, at the cost of higher mortality during the epidemic.

Under our baseline calibration, a comparison of the utilitarian optimal policy to the actual policy in place as of Easter 2020 indicates that the shutdown in place was around twice as extensive as it should be. However, the optimal policy calls for leaving a partial shutdown in place well into the fall. Ending the shutdown at Easter as originally contemplated by the current U.S. administration would have implied an additional 231,000 deaths. Ending the shutdown at the end of June is predicted to lead to a second wave of infections, in line with the current empirical record.

We also ask how preferred policies change if a vaccine is expected by the end of 2020, an optimistic time frame. We find that individuals then prefer more extensive mitigation for a longer period than in our baseline economy. Without a vaccine, economic mitigation effectively delays the total number of infections, but some of the people who avoid infection early on simply become infected later. Since the vaccine immediately stops the spread of the virus, people are more willing to give up consumption in exchange for eliminating the chance of becoming infected in the future.

A second extension we explore is universal antibody testing, which allows the planner to avoid unnecessarily mitigating workers who have recovered from the disease, and can neither become sick again nor infect others. Perhaps surprisingly, we find this is not a very valuable policy tool.

Third, we explore one intervention designed to explicitly protect the old by insulating them from infection risk while shopping. Given this sharper tool for reducing old-age mortality risk, the planner optimally chooses to reduce the extent of mitigation, thereby moderating the decline in aggregate output and consumption.

Our paper contributes to a rapidly expanding literature on the interaction between health pandemics and economic activity, with primary focus on the current COVID-19 crisis. Atkeson (2020) was perhaps the first to introduce economists to the epidemiological SIR class of models. He emphasizes the negative outcomes that arise if and when the fraction of active infections in the population exceeds 1% (at which point the health system is predicted to be severely challenged) and 10% (which may result in severe staffing shortages in key financial and economic infrastructure sectors) as well as the cumulative burden of the disease over an 18-month horizon. Fernandez-Villaverde and Jones (2020) use high-frequency data to estimate the SIR model for a rich set of countries and cities, and use the model to provide forward simulations of potential paths the epidemic can take, paying close attention to the importance of time-varying parameters as well as parameter uncertainty in the model. Greenstone and Nigam (2020) use the Imperial College epidemiological model (Flaxman et al. 2020) to compare the paths under moderate social distancing versus no policy action and use the statistical value-of-life approach to assess

the social cost of no action. They calculate 1.7 million lives saved between March 1 and October 1 from social distancing, 37% of them due to less overcrowding in hospitals.

Eichenbaum et al. (2020) extend the canonical SIR 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 effects 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 much smaller if individuals can endogenously adjust the sectors in which they consume. Argente et al. (2020) characterize the optimal path of mitigation policies in a representative agent economy, and Acemoglu et al. (2020) as well as Brotherhood et al. (2020) study optimal age-based policies. Relative to these papers our work stresses the large heterogeneity in the costs and benefits of mitigation policies and the importance of the cost of redistribution for the extent of optimal mitigation.

Toxvaerd (2020) and Farboodi et al. (2020) characterize the simultaneous determination of infections and social distancing in models where individuals make conscious choices about their social interactions. Moll et al. (2020) develop 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. Guerrieri et al. (2020) develop a theory of Keynesian supply shocks that induce changes in aggregate demand which amplify the adverse economic consequences of the original pandemic-induced supply shocks. 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.

An important literature explores the potential role of testing (and tracing). Berger et al. (2020) and Chari et al. (2020) extend the baseline Susceptible-Exposed-Infectious-Recovered (SEIR) infectious-disease model to explore the role of testing for the implementation of selective social separation policies. Using the Chinese experience, Fang et al. (2020) quantify the causal impact of human mobility restrictions and find that the lock-down was very effective, providing estimates of diffusion under different scenarios. Hall et al. (2020) provide a set of illuminating calculations to assess the value of never having had the virus. They find that people would be willing to pay about a quarter of one year's worth of consumption.

In Section 2, we start by describing how we model the joint evolution of the economy and the population. In Section 3, we then turn to describe how we model mitigation and redistribution policies and how we go about solving for optimal policies. The calibration strategy is described in Section 3.3. The findings are in Section 4, and Section 5 discusses extensions to vaccination,

antibody testing and additional protection of the old. Section 6 concludes.

2 The Model

We first describe the individual state space, setting out the nature of heterogeneity by age and health status. In Section 2.2, we describe the multi-sector production technology and how *mitigation* shapes the pattern of production. Section 2.3 explains the details of our *SAFER* extension of the standard *SIR* epidemiological model and the channels of disease transmission.

2.1 Household Heterogeneity

Agents can be young or old, denoted by y and o. We think of the young as below the age of 65 and they will comprise $\mu^y = 85$ percent of the population. For simplicity, and given the short time horizon of interest, we abstract from population growth and ignore aging. Within each age group, agents are differentiated by health status, which can take six different values: susceptible s, asymptomatic a, miserable with a fever 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 – they are subsets of the infected group in the standard *SIR* model – and can pass it onto others. However, they differ in their symptoms. The asymptomatic have no symptoms or very mild ones and thus unknowingly spread the virus. We model this state explicitly (in contrast to the prototypical *SIR* model) because a significant percentage of individuals infected with COVID-19 experience no symptoms.[1] Those with a fever are sufficiently sick to know they are likely contagious, and they stay at home and avoid the workplace and market consumption. Those requiring emergency care are hospitalized. The recovered are again healthy, no longer contagious, and immune from future infection. A worst-case virus progression is from susceptible to asymptomatic to fever to emergency care to dead.[2] However, recovery is possible from the asymptomatic, fever, and emergency-care states.

2.2 Activity: Technology and Mitigation

Young agents in the model are further differentiated by the sector in which they can work. A fraction μ^b of the young work in the basic (essential) sector, denoted b, while the rest, $1 - \mu^b$, work in a luxury (non-essential) sector, denoted ℓ. We assume that output of the basic sector is so vital that it is never optimal to send home even a subset of b sector workers. 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. We will call such a policy a

---

[1]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.

[2]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 state is a hybrid of 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 the SEIR I state). Berger et al. (2020) make a similar modeling choice.

(macroeconomic) mitigation policy, m. More precisely, m_t will denote the fraction of luxury workers that are instructed to 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 is no scope for confusion. Technology parameters are denoted with Greek letters. Generically, we use the letter x to denote population measures, with superscripts specifying subsets of the population. These super-indices index age, sector, and health status, in that order. For example, x^{ybs} is the measure of young individuals working in the basic sector who are susceptible.

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:

$$y^b = x^{ybs} + x^{yba} + x^{ybr}. \tag{1}$$

Note that this specification assumes that those asymptomatic individuals carrying the virus continue to work.[3] In contrast, we assume those with fever stay at home. Output in the luxury sector, in contrast, does depend on the mitigation policy and is given by

$$y^\ell = (1 - m_t) \left(x^{y\ell s} + x^{y\ell a} + x^{y\ell r} \right). \tag{2}$$

We assume that both sectors produce the same good and are perfect substitutes.[4] Under this assumption, total output of the single consumption good is determined by

$$y = y^b + y^\ell. \tag{3}$$

We assume that a fixed amount of output $\eta\Theta$ is spent on emergency hospital care, where Θ is the capacity of hospital beds, and η is the cost of providing and maintaining one bed.

In practice, different sectors of the economy are heterogeneous with respect to the extent to which production and consumption generate risky social interaction. For example, some types of work and market consumption can easily be done at home, while for others, avoiding interaction is much harder. A sensible shutdown policy will first shutter those sub-sectors of the luxury sector that generate the most interaction. Absent detailed micro data on social interaction by sector, we model this in the following simple way.[5] Assume workers are assigned to a unit interval

---

[3]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.

[4]We make this assumption primarily for the sake of tractability. If outputs of the two sectors were complementary, there would be changes in relative prices and wages when output of the luxury sector was suppressed.

[5]See Xu et al. (2020) for more detailed evidence on infection patterns in the workplace.

of sub-sectors $i \in [0, 1]$ where sub-sectors are ranked from those generating the least to those generating the most social interaction. Also assume the sector-specific infection-generating rates are $\beta_w^i = 2\alpha_w i$ and $\beta_c^i = 2\alpha_c i$, where (α_w, α_c) are parameters, to be calibrated below, governing the intensity by which meetings among individuals generate infections. When the government asks fraction m_t of luxury workers to stay at home, assume it targets the sub-sectors generating the most interactions, that is, $i \in [1 - m_t, 1]$. The average interaction rates of the sectors that remain are then $\alpha_w(1 - m_t)$ and $\alpha_c(1 - m_t)$, respectively.[6] Because the government cannot shut down any basic sub-sectors of the economy, the economy-wide work-related infection-generating probability is then given by

$$\beta_w(m_t) = \frac{y^b}{y(m_t)}\alpha_w + \frac{y^\ell(m_t)}{y(m_t)}\alpha_w(1 - m_t),$$

with an analogous expression for $\beta_c(m_t)$. The key property of this expression is that as mitigation is increased, the average social interaction-generating rate will fall.

2.3   Health Transitions: The SAFER Model

We now describe the dynamics of individuals across health states. 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 interest of compact notation, we will 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 $x = \sum_{i \in \{s,a,f,e,r\}} x^i = x^{yb} + x^{y\ell} + x^o$ denote the entire living population.

In our model, the crucial health transitions that can be affected by mitigation policies are from the susceptible to the asymptomatic state. Equations (4)-(6) below capture the flow of basic sector workers, luxury sector workers, and older individuals out of the susceptible state and into the asymptomatic state. The number of such workers who catch the virus is their original mass (x^{ybs} for young basic sector workers, for example) times the number of virus-transmitting interactions they have (the term in square brackets). We model four sources of possible virus contagion: people can catch the virus from colleagues at work, from market consumption activities, from family or friends outside work, and from taking care of the sick in hospitals. The four terms in the bracket capture these four sources of infection, which we index w, c, h, and e, respectively. For a given type of individual, the flow of new infections from each of these activities is the product of the number of contagious people they can expect to meet, which we denote $\mu_j(m_t)$ for $j \in \{w, c, h, e\}$, and the likelihood that such meetings result in infection, which is the infection-generating rate described above, $\beta_j(m_t)$. For work and consumption activities, both the number of contagious people in a given setting and the rate at

---

[6] $E\left[\alpha_w i | i \leq (1 - m_t)\right] = \frac{2\alpha_w}{1 - m_t} \int_0^{1 - m_t} i \, di = \frac{2\alpha_w}{1 - m_t} \frac{(1 - m_t)^2}{2} = \alpha_w(1 - m_t)$.

which they transmit the virus potentially depend on the level of economic mitigation m_t.

$$\dot{x}^{ybs} = -\left[\beta_w(m_t)\mu_w(m_t) + \beta_c(m_t)\mu_c(m_t) + \beta_h\mu_h + \beta_e\mu_e\right] x^{ybs} \tag{4}$$

$$\dot{x}^{y\ell s} = -\left[\beta_w(m_t)\mu_w(m_t)(1 - m_t) + \beta_c(m_t)\mu_c(m_t) + \beta_h\mu_h\right] x^{y\ell s} \tag{5}$$

$$\dot{x}^{os} = -\left[\beta_c(m_t)\mu_c(m_t) + \beta_h\mu_h\right] x^{os} \tag{6}$$

where the relevant population shares in the above expressions are given by

$$\mu_w(m_t) = x^{yba} + (1 - m_t)x^{y\ell a} \tag{7}$$

$$\mu_c(m_t) = x^a y(m_t) \tag{8}$$

$$\mu_h = x^a + x^f \tag{9}$$

$$\mu_e = x^e \tag{10}$$

Consider the first outflow rate in equation (4). The flow of young basic sector workers getting infected at work, $\beta_w(m_t)\mu_w(m_t)$, is the probability of a virus-spreading interaction per contagious worker, $\beta_w(m_t)$, times the number of contagious workers, which is defined in equation (7). Note that we are assuming that people with symptoms always stay at home (a minimal precaution) and that basic and luxury workers mingle together at work.

The flow $\beta_c(m_t)\mu_c(m_t)$ of young basic sector workers getting infected from market consumption is similarly constructed. We assume that the number of consumption-related infections is proportional to the number of asymptomatic individuals in the population and to the level of economic activity, which is identical to the number of workers (see equation 8).[7] Note that we are assuming that people with symptoms stay at home and do not go shopping.

The rate at which a young basic worker contracts the virus at home, $\beta_h\mu_h$, depends on the number of contagious workers in the household, μ_h defined in equation (9). Note that both asymptomatic and fever-suffering workers reside at home. 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, $\mu_e = x^e$, with infection-generating rate β_e, which reflects the strength of precautions taken in hospitals.

Parallel to equation (4), equation (5) describes infections for the susceptible population working in the luxury sector. For this group, the risks of infection from market consumption and at home are identical to those for basic sector workers. However, individuals in this sector

[7]Note that we have assumed that the number of shopping-related infections for a given type is proportional to economy-wide output, rather than to the type-specific level of consumption. One interpretation of this assumption is that each consumer visits each store in the economy and faces a similar infection risk irrespective of how much they spend. The common infection risk is proportional to the equilibrium number of stores, which in turn is proportional to the aggregate employment level.

work reduced hours when $m_t > 0$ 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 (4) is absent in equation (5). Similar to equation (4) and equation (5), equation (6) displays infections among the old. They get infected only from market consumption and from interactions at home.

The remainder of the epidemiological block is fairly mechanical and simply describes the transition of individuals though the health states (asymptomatic, fever-suffering, hospitalized, and recovered) once they have been infected. These transitions are described in equations equation (31) to equation (42) in Appendix A, with parameters that are allowed to vary by age. Transition into death occurs from the emergency care state at age-dependent rates $\sigma^{yed} + \varphi$ and $\sigma^{oed} + \varphi$, where φ is the excess mortality rate when hospital capacity is overused, and is given by

$$\varphi = \lambda_o \max\{x^e - \Theta, 0\}. \tag{11}$$

In equation 11 the term in the max operator defines the extent of hospital overuse given capacity Θ, which we treat as fixed in the time horizon analyzed in this paper.[8] The parameter λ_o controls how much the death rate of the hospitalized rises (and the recovery rate falls) once capacity is exceeded.

2.4 Preferences

Preferences incorporate utility from both being alive and being in a specific health state. Lifetime utility for the old is given by

$$E\left\{\int e^{-\rho_o t} \left[u(c_t^o) + \bar{u} + \widehat{u}_t^j\right] dt\right\}, \tag{12}$$

where expectations are taken with respect to the random timing of death, and where $\bar{u}$ measures the flow utility from being alive (the utility of being dead is implicitly zero). Similarly, $\widehat{u}_t^j$ is the intrinsic utility of being in state health j. We will assume that $\widehat{u}_t^s = \widehat{u}_t^a = \widehat{u}_t^r = 0$, while $\widehat{u}_t^e < \widehat{u}_t^f < 0$. Thus, having a fever 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 $u(c_t^o)$ and discount the future at rate ρ_o.

Symmetrically, the young also care about their consumption c_t^y, as well as about their health

---

[8]When solving for the non-parametric optimal mitigation policy in Section 4.3, we use the smooth approximation

$$\max\{x^e - \Theta, 0\} \approx \frac{\log\left(1 + e^{N(x^e - \Theta)}\right)}{N}.$$

The approximation error is always less than 0.04% of peak hospitalizations with $N = 1000000$.

and about being alive, according to the lifetime utility function:

$$E\left\{\int e^{-\rho_y t}\left[u\left(c_t^y\right)+\bar{u}+\widehat{u}_t^j\right]dt\right\}.\qquad(13)$$

In our calibration, we will impose $\rho_o > \rho_y$ as a simple way to capture higher life expectancy for the young. As a result, while young and old enjoy the same flow value from being alive, the present value of this value will be lower for the old.

Note that workers who have a fever or are in the hospital do not work. Neither does a fraction m of luxury sector workers whose workplaces have been shut down by mitigation policy. Therefore, in equilibrium, young workers will experience different consumption depending on whether they work. Thus, the expected utility of a worker will depend for two reasons on the sector in which she works. First, sectors differ in the share of economic activity being shut down (and thus, for the individual worker, in the probability of being able to work when healthy). Second, a worker's sector will affect her distribution of health outcomes.[9]

3 The Public Sector

In Section 3.1 we describe the government policy tools, 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, 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 from workers to individuals who do not or cannot work: those unemployed because of shutdowns, those with fever or hospitalized, and those who have retired. All workers share a common consumption level c^w and all individuals not working share a common consumption level c^n.[10] The redistribution policy choice is how much to transfer, in each instant t, from the working to the non-working 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 of the health population distribution x and current mitigation

---

[9]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.

[10]This is the allocation chosen by a government that equally values all individuals (equal Pareto weights). It is also the only allocation that is feasible if the government can observe an individual's income but not her sector, age, or health status.

$m = m_t$.[11] These are defined as

$$\mu^n(m, x) = x^{y\ell f} + x^{y\ell e} + x^{ybf} + x^{ybe} + m\left(x^{y\ell s} + x^{y\ell a} + x^{y\ell r}\right) + x^o, \tag{14}$$

$$\mu^w(m, x) = x^{ybs} + x^{yba} + x^{ybr} + [1 - m]\left(x^{y\ell s} + x^{y\ell a} + x^{y\ell r}\right), \tag{15}$$

$$v^w(m, x) = \frac{\mu^w(m, x)}{\mu^w(m, x) + \mu^n(m, x)}, \tag{16}$$

where $v^w(m, x)$ is the share of working individuals in the population. The aggregate resource constraint can then be written as

$$\mu^w c^w + \mu^n c^n + \mu^n T(c^n) = y - \eta\Theta = \mu^w - \eta\Theta \tag{17}$$

where $y = \mu^w$ since each working individual produces one unit of output.

Notice that there are no dynamic consequences of the transfer choice c^n. In particular, this choice has no impact on any health transitions. At each date t, we can therefore solve a static optimal transfer problem (given the current level of mitigation $m = m_t$) that delivers a maximum level of instantaneous 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 for the optimal mitigation problem of the government. Assuming that all individuals have log-utility and receive the same social welfare weights, the function $W(x, m)$, is given by

$$W(x, m) = \max_{c^n, c^w} \left[\mu^w \log(c^w) + \mu^n \log(c^n)\right] + (\mu^w + \mu^n)\bar{u} + \sum_{i,j} x^{ij}\hat{u}^j, \tag{18}$$

for $i \in \{yb, y\ell, o\}$ and $j \in \{s, a, f, e, r\}$, where the maximization is subject to the aggregate resource constraint (17). Defining net per-capita income $\tilde{y}$ and average transfer costs $t(c^n)$ as

$$\tilde{y} = v^w - \frac{\eta\Theta}{\mu^w + \mu^n} \tag{19}$$

$$t(c^n) = \frac{T(c^n)}{c^n}, \tag{20}$$

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

$$v^w c^w + (1 - v^w)c^n(1 + t(c^n)) = \tilde{y}. \tag{21}$$

---

[11]We will suppress the dependence on (x, m) when there is no room for confusion.

The following results then follow directly from the first order condition when maximizing (18) subject to (21).

Proposition 1 *The optimal solution to the government transfer problem is given by the solution to the following system:*

$$\frac{c^w}{c^n} = 1 + T'(c^n) \tag{22}$$

$$c^n \left[1 + v^w T'(c^n) + (1 - v^w)t(c^n)\right] = \tilde{y} \tag{23}$$

Period welfare is given by

$$W(x, m) = \left[\mu^w + \mu^n\right] w(x, m) \tag{24}$$

$$w(x, m) = \log(c^n) + v^w \log(1 + T'(c^n)) + \bar{u} + \sum_{i,j} \frac{x^{ij}}{\mu^w + \mu^n} \widehat{u}^j, \tag{25}$$

where the only endogenous input in the period welfare function is c^n which is tied to per capita output and thus mitigation via equation (23). Given the allocations (c^w, c^n), we can also construct period welfare for each household type $\{yb, y\ell, o\}$. We do so in Appendix B.

Note that since $\mu^w + \mu^n$ is independent of mitigation, we can discuss the impact of mitigation on current welfare in terms of the per capita welfare function $w(x, m)$. First, mitigation lowers per capita income and, through it, the level of consumption c^n. Second, the transfer cost to non-working households distorts risk sharing. In particular, when the marginal cost of transfers is positive, there is a wedge between consumption of workers and non-workers. Thus an increase in mitigation increases consumption inequality, making mitigation less attractive (an increase in m reduces v^w and thus the positive welfare contribution from the second term in eq. (25)). To see these effects most clearly, consider two special cases.

Corollary 1 *Assume that the transfer cost is linear such that $T(c^n) = \tau c^n$. In this case the optimal allocation is given by:*

$$c^w = \tilde{y} \quad and \quad c^n = \frac{\tilde{y}}{1 + \tau}$$

$$w(x, m) = \log(\tilde{y}) - (1 - v^w)\log(1 + \tau) + \bar{u} + \sum_{i,j} \frac{x^{ij}}{\mu^w + \mu^w} \widehat{u}^j.$$

The negative economic impact of mitigation is given by

$$\frac{\partial w(x, m)}{\partial m} = \frac{\partial \tilde{y}}{\partial m} + (1 + \tau)\frac{\partial v^w}{\partial m} < 0, \tag{26}$$

since both $\frac{\partial \tilde{y}}{\partial m}$ and $\frac{\partial v^w}{\partial m}$ are negative. In addition, the larger is the marginal cost of transfers τ, the more negative is $(1 + \tau)\frac{\partial v^w}{\partial m}$ and thus $\frac{\partial w(x,m)}{\partial m}$.

The first term in (26) is the direct impact of mitigation on output, while the second captures the negative effect on inequality. Note how mitigation and redistribution costs interact: the larger is the marginal cost of redistribution τ, 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[12] $T(c^n) = \frac{\tau}{2}\frac{\mu^n}{\mu^w}(c^n)^2 = \frac{\tau}{2}\left(\frac{1-v^w}{v^w}\right)(c^n)^2$.

Corollary 2 *For the quadratic specification, the optimal allocations are given by:*

$$c^n = \frac{\sqrt{1 + 2\tau\frac{1-(v^w)^2}{v^w}\tilde{y}} - 1}{\tau\frac{1-(v^w)^2}{v^w}} \tag{27}$$

$$c^w = c^n(1 + T'(c^n)) = c^n\left(1 + \tau\frac{1-v^w}{v^w}c^n\right). \tag{28}$$

The term $\left(1 + \tau\frac{1-v^w}{v^w}c^n\right)$ 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 v^w and increases the share of non-workers $1 - v^w$, the effective price of transfers at the margin increases with mitigation, and the price rises faster with mitigation the higher is τ.

3.3 Optimal Policy

We now assume there is a government/planner (we use these names synonymously, 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

$$\max_{m(t)} \int_0^\infty e^{-\rho t}W(x)\ dt, \tag{29}$$

subject to the laws of motion of the health population distribution.

The optimal policy path is the solution to this optimal control problem. We formally state that problem in Appendix C. The key trade-off with mitigation efforts m is that a marginal

---

[12] Then 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 $\frac{\tau}{2}\left(\frac{\mu^n c^n}{\mu^w}\right)^2$. Multiplying this by the total number of workers μ^w gives the total transfer cost. The quadratic form is chosen for analytical convenience but is not central for our qualitative arguments.

increase in m entails static economic costs of $\frac{\partial W(x,m)}{\partial m}$ stemming from reduced output and more consumption inequality. The dynamic benefit is a favorable change in the population health distribution: an increase in m reduces the outflow of individuals from the susceptible to the asymptomatic state.

We will start our optimal policy analysis by approximating the optimal time path of mitigation by functions that are part of the following parametric class of generalized logistic functions of time:

$$m(t) = \frac{\gamma_0}{1 + \exp(-\gamma_1(t - \gamma_2))}. \tag{30}$$

The advantage of this class of functions is that its parameters, and thus the optimal mitigation path itself, are readily interpretable. Specifically, the parameter γ_0 controls the level of mitigation at $t = 0$. The parameter γ_2 governs when mitigation is reduced, and the parameter γ_1 commands the swiftness with which mitigation is reduced. Note that as $t \to \infty$, $m(t) \to 0$. In Section 4.3 we will solve the unconstrained optimal control problem and show that the additional welfare gains, relative to the parametric optimal path, are rather small. very

Our calibration procedure has two parts. The first involves selecting a large set of parameter values in a standard way based on a mix of external evidence and choices about empirical counterparts to model objects. The second part is more delicate, and has to do with quantifying the aggregate evolution of the pandemic in its early stages. In this second step we need to estimate changes in behavior and policies as the United States moved from business-as-usual to a partial lock-down coupled with a set of behavioral changes designed to reduce the spread of infection. We begin with the first, and more straightforward, part of the calibration.

We set the population share of the young, μ^y, to be 85%, which is the current fraction of the US population below the age of 65.

Preferences   We assume logarithmic utility from consumption:

$$u(c) = \log c.$$

We set the pure time discount rate in annual terms to 3%. To accommodate differential mortality by age in the simplest way we assume that 500 days after the start of the pandemic (sufficient for it to have run its course), the discount rate becomes 4% for the young and 10% for the old. These values are chosen to reflect, respectively, a residual expected duration of life of 47.5 years for a 32.5 year old, and of 14 years for a 72.5 year old, numbers which are consistent with recent pre-COVID-19 life tables.

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 VSL of $11.5 million (see Greenstone and Nigam 2020). This is a high value, relative to values used in other contexts. Assuming an average of 37 residual life years discounted at a 3% rate, this translates

into 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 statistical life calculation,

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

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 $\bar{u} = 11.4 - \ln \bar{c}$. This implies an easily interpretable trade-off between mortality risk and consumption. For example, 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 for the agent 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%.

For the disutility of fever, we define $\hat{u}^f = -0.3 \left(\ln(\bar{c}) + \bar{u} \right)$, following Hong et al. (2018). We set $\hat{u}^e = - \left(\ln(\bar{c}) + \bar{u} \right)$, so that the flow value of being in hospital is equal to the flow value of being dead (zero).

Sectors   To calibrate the employment and output share of the basic sector of the economy, μ^b, we use BLS employment shares by industry. We categorize the following industries as basic: agriculture, health care, financial activities, utilities, and federal government. Mining, construction, manufacturing, education, and leisure and hospitality are allocated to the luxury sector. The remaining industries are assumed to be a representative mix of basic and luxury. This partition implies that pre-COVID, the basic sector accounts for $\mu^b = 45.4$ percent of the economy.

Redistribution   We adopt the quadratic formulation of transfer costs described above. 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 for taxpayers of around $1.38 (Saez et al. 2012). This suggests $\tau \frac{1-v^w}{v^w} c^n = 0.38$. Given the first order condition above, this means that an optimal redistribution scheme would imply $c_n/c_w = 1/1.38 = 0.72$ in pre-COVID times. Moreover, given $\eta\Theta = 0.021$, $\tau \frac{1-v^w}{v^w} c^n = 0.38$, and $v = \mu^y = 0.85$, corollary 2 implies $\tau = 3.51$.

Hospital Capacity   Tsai et al. (2020) estimate that 58,000 ICU beds are potentially available nationwide to treat COVID-19 patients. However, only 21.5% of COVID-19 hospital admissions require intensive care, suggesting that total hospital capacity is around $58,000/0.215 = 270,000$. Tsai et al. (2020) emphasize that this capacity is very unevenly allocated geographi-

cally, and in addition, there is significant geographic variation in virus spread. Thus, capacity constraints are likely to bind in more and more locations as the virus spreads. We therefore set $\Theta = 100,000$, so that hospital mortality starts to rise when 0.042 percent of the population is hospitalized. Because the cost of a day in intensive care is around \$7,500, we set $\eta = 50$, so emergency care consumes about 2.1 percent of pre-COVID output.[13] We set the parameter λ_o such that the mortality rate in emergency care for the old at the peak of the epidemic in our simulation in which economic mitigation ends on April 21 is 20 percent above its value when capacity is not exceeded.[14]

Disease Progression   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 worst health status and the chance of recovery at the three infectious stages: asymptomatic, feverish, and hospitalized. We assume that young and old exit each stage at the same rate but potentially differ in 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 fever and 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, fever, 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, respectively, with these durations common across age groups. The exit rate from the asymptomatic state to recovery defines the number of asymptomatic cases of COVID-19 and is an important but uncertain parameter. We assume that asymptomatic recovery and progression to the fever-suffering state are equally likely.[15]

We let the relative rates of recovery from the fever and emergency care states vary with age, to reflect the fact that infections in older individuals are much more likely to require hospitalization and hospitalizations are also somewhat more likely to lead to death. We set the recovery rate from fever to 96% for the young and 75% for the old, based on evidence from Table 1 of the Imperial College study (Ferguson et al. 2020). Similarly, given evidence on differential mortality rates, we set the recovery rates from the emergency care state to 95% for the young and 80% for the old (assuming no hospital overuse). Given these choices, the probability that a newly infected young individual will ultimately die from COVID-19 is $0.5 \times 0.04 \times 0.05 = 0.1\%$,

---

[13]Total health care spending in the United States is 18% of GDP. Of this, around one third is spending on hospitals.

[14]Much of the concern about exceeding capacity has focused on a potential shortage of ventilators. However, recent evidence from New York City indicates that 80% of ventilated COVID-19 patients die, suggesting a limited maximum potential excess mortality rate associated with this particular channel.

[15]Given that the asymptomatic state has roughly half the duration of the fever 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.

while the conditional probability, conditional on developing fever, is 0.2%. The corresponding numbers for an older individual are 2.5% and 5.0%.

Sources of Infection   Given the σ parameters, the parameters α_w, α_c, β_h, and β_e determine the rate at which contagion grows over time. We set β_e, the hospital infection-generating rate, so that this channel accounts for 5 percent of cumulative COVID-19 infections though April 12. This implies $\beta_e = 0.80$.[16] The values of α_w, α_c, 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, via market consumption, and in non-market settings.

Mossong et al. (2008) find that 35% of potentially infectious inter-person contact happens in workplaces and schools, 19% occurs in travel and leisure activities, and the remainder is in home and other settings.[17] These shares should be interpreted as reflecting behavior in a normal period of time, rather than in the midst of a pandemic. We associate workplace and school transmission with transmission at work, travel and leisure with consumption-related transmission, and the residual categories with transmission at home. These targets are used to pin down choices for α_w and α_c, both relative to 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 and zero mitigation ($m = 0$), R_0^y is given by

$$R_0^y = \frac{\alpha_w x^y + \alpha_c \mu^y + \beta_h}{\sigma^{yar} + \sigma^{yaf}} + \frac{\sigma^{yaf}}{\sigma^{yaf} + \sigma^{yar}} \frac{\beta_h}{\sigma^{yfr} + \sigma^{yfe}} + \frac{\sigma^{yaf}}{\sigma^{yaf} + \sigma^{yar}} \frac{\sigma^{yfe}}{\sigma^{yfe} + \sigma^{yfr}} \frac{\beta_e x^{yb}}{\sigma^{yer} + \sigma^{yed}}$$

where this expression exploits the fact that when $m = 0$, $\beta_w(0) = \alpha_w$ and $\beta_c(0) = \alpha_c$.

The logic is that this individual will spread the virus while asymptomatic, suffering fever, and hospitalized —the three terms in the expression. They expect to be asymptomatic for $(\sigma^{yar} + \sigma^{yaf})^{-1}$ days, feverish (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 fever state is $\frac{\sigma^{yaf}}{\sigma^{yaf} + \sigma^{yar}}$, and the chance they reach the emergency room is the product $\frac{\sigma^{yaf}}{\sigma^{yaf} + \sigma^{yar}} \frac{\sigma^{yfe}}{\sigma^{yfe} + \sigma^{yfr}}$. While asymptomatic, they spread the virus both at work and at home, and pass the virus on to $\alpha_w x^y + \alpha_c \mu^y + \beta_h$ susceptible individuals per day.[18] While feverish, they stay at home and pass the virus to β_h individuals per day. While sick they pass it to $\beta_e x^{yb}$ basic

---

[16]5% is an estimate by Kent Sepkowitz (Memorial Sloan Kettering Cancer Center) of the share of infections accruing to health-care workers who acquired the infection after occupational exposure: https://www.cnn.com/2020/04/15/opinions/health-care-deaths-sepkowitz-opinion/index.html

As of March 24th, 14% of Spain's confirmed cases were health care workers: https://www.nytimes.com/2020/03/24/world/europe/coronavirus-europe-covid-19.html

[17]Xu et al. (2020) discuss in detail heterogeneity in contact rates across different types of business (closed office, open office, manufacturing and retail) and a range of interventions that can reduce those rates.

[18]Recall that x^y is the pre-COVID number of workers, and α_w is the probability that transmission occurs when an infected worker meets a susceptible one. Recall that we assume consumption contagion is proportional to output, and pre-COVID output is $\mu^y = x^y$.

workers per day in hospital.

The reproduction number expression for an old asymptomatic person, R_0^o, is similar, except that the old do not pass the virus on at work, but are more likely to require hospitalization and transmit the virus in hospital.[19] For the population as a whole, the overall R_0 is a weighted average of these two group-specific values: $R_0 = \mu^y R_0^y + (1 - \mu^y) R_0^o$.

The expected shares of workplace and consumption transmission are given by

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

$$\frac{\text{consumption transmission}}{\text{all transmission}} = \frac{1}{R_0} \left[\mu^y \left(\frac{\mu^y \alpha_c}{\sigma^{yar} + \sigma^{yaf}} \right) + (1 - \mu^y) \left(\frac{\mu^y \alpha_c}{\sigma^{oar} + \sigma^{oaf}} \right) \right].$$

Given these equations, we set the relative values α_w/β_h, α_c/β_h to replicate shares of workplace and consumption transmission equal to 35% and 19%. Note that this evidence does not pin down the *levels* of α_w, α_c, and β_h, to which we now turn.

Table 1: Epidemiological Parameter Values

	Behavior-Contagion		
α_w	infection at work	35% of infections	0.30
α_c	infection through consumption	19% of infections	0.15
β_e	infection in hospitals	5% of infections at peak	0.80
β_h	infection at home	remaining infections	0.18
ζ	scale of social distancing	deaths on April 12	0.38
R_0 (pre 3/21)	effective virus reproduction	composite parameter	3.61
R_0 (3/21 to 4/12)	effective virus reproduction	composite parameter	1.02
	Disease Evolution		
σ^{yaf}	rate for young asymptomatic into fever	50% fever, 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 fever	50% fever, 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 fever into emergency	4% hospitalization, 10 days	$\frac{0.04}{10}$
σ^{yfr}	rate for young fever into recovered		$\frac{0.96}{10}$
σ^{ofe}	rate for old fever into emergency	25% hospitalization, 10 days	$\frac{0.25}{10}$
σ^{ofr}	rate for old fever into recovered		$\frac{0.75}{10}$
σ^{yed}	rate for young emergency into dead	0.2% 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		$\frac{0.80}{8}$

History, R_0, and Initial Conditions   We will think of a policy maker choosing a path for mitigation m_t starting from April 12, 2020. The dynamics of the disease going forward, and thus the optimal path for m_t, will be highly sensitive to the distribution of the population

---

[19] $R_0^o = \frac{\alpha_c \mu^y + \beta_h}{\sigma^{oar} + \sigma^{oaf}} + \frac{\sigma^{oaf}}{\sigma^{oaf} + \sigma^{oar}} \frac{\beta_h}{\sigma^{ofr} + \sigma^{ofe}} + \frac{\sigma^{oaf}}{\sigma^{oaf} + \sigma^{oar}} \frac{\sigma^{ofe}}{\sigma^{ofe} + \sigma^{ofr}} \frac{\beta_e x^{yb}}{\sigma^{oer} + \sigma^{oed}}.$

by health status at this date. Given limited testing, we have limited information about this distribution.

In addition, going forward, the dynamics of the disease will depend on the basic reproduction number R_0, which in our model is determined at a structural level by the levels of the infection-generating parameters α_w, α_c, and β_h. Existing estimates for R_0 for COVID-19, absent additional social distancing measures or economic shutdowns, are in the range of 2 to 4 (e.g., Flaxman et al. 2020). But given all the precautions that Americans are currently choosing to take or being required to take, the current effective R_0 is likely much lower.

To pin down the April 12 health status distribution and the April 12 level for the infection-generating parameters, we take the following approach. First, we will assume that America changed on March 21. Before that date, people behaved as normal, and none of the economy was shuttered, corresponding to $m = 0$. On March 21, we assume infection-generating rates fell discretely and proportionately to new lower levels $\zeta \alpha_w$, $\zeta \alpha_c$, and $\zeta \beta_h$ with $\zeta < 1$ (we assume no change in the hospital infection-generating rate β_e). In addition, and at the same date, we assume that states introduced measures that effectively shut down a fraction $m = 0.5$ of the luxury sector, therefore immediately idling $0.5(1 - \mu^b) = 27.5$ percent of the workforce. Bick and Blandin 2020 report that from February 9-15 to April 12-18, weekly hours worked per working age adult declined 27 percent.[20] In reality, changes in social distancing practices and shutdowns happened more gradually, but March 21 seems a natural focal date: California announced the closure of non-essential businesses on March 19, and New York and Illinois did so on March 20.

Of the data we have on health outcomes, the most reliable are for the number of deaths attributable to COVID-19. We will therefore target three specific moments involving deaths: the cumulative number of deaths up to March 21 (343), the cumulative number as of April 12 (22,055), and the three-day moving-average number of deaths per day on April 12 (1,632).[21] To hit these target moments, we treat as free parameters (1) β_h —the pre-March 21 infection-generating rate at home; (2) ζ, the proportional amount by which infection-generating rates fall on March 21; and (3) the initial number of infections at the date we start our model simulation, which is February 12.

To understand how this identification scheme works, consider that the death toll rose from 343 to 22,055 deaths in only three weeks, but the number of daily deaths was not especially high (nor was it growing especially fast) at the end of this period. This suggests that there were already many infections in the pipeline on March 21, but those infections did not grow rapidly from March 21 onward, which indicates a low value for ζ. At the same time, a high level of

---

[20]Cajner et al. (2020) estimate that private-sector employment fell 21 percent between mid February and late April.

[21]Death tallies vary slightly across data sources and are occasionally revised retrospectively. We use the New York Times numbers: https://www.nytimes.com/interactive/2020/us/coronavirus-us-cases.html?action=click&module=Spotlight&pgtype=Homepage

March 21 infections is informative about the level of initial infections on February 12. Finally, a large number of infections on March 21, but low death toll up to that date, points to high R_0 (and a high β_h) before March 21: rapid spread can deliver lots of new infections without many deaths (yet).

This calibration strategy yields an initial effective R_0 before March 21 of 3.61, which falls to 1.02 after March 21, reflecting a value for ζ of 0.38. Part of this decline reflects the start of economic mitigation. Absent mitigation (with $m = 0$), the effective R_0 after March 21 would be 1.44. This calibration implies the distribution of the population by health status summarized in Table 3. Thus, the calibration implies that 1.52% of the US population was actively infected (including asymptomatic infections) on March 21, with that number rising to 1.75% by April 12, with an additional 3.9% having recovered.[22]

For the time path of mitigation, our baseline simulation, designed to approximate current US policy, will assume $m = 0.5$ for 100 days from March 21 onward, followed by $m = 0$ thereafter. Thus, extensive economic shutdowns are in place until June 29, and then shutdowns are abruptly ended. This path is implemented in the context of the mitigation function (eq. 30) by setting $\gamma_0 = 0.5$, $\gamma_1 = -0.5$, and $\gamma_2 = 100$. Note that we assume no change in the infection-generating parameters α_w, α_c or β_h from March 21 onwards; thus relaxing economic mitigation does not imply an end to all social distancing. All the epidemiological and economic parameter values are summarized in Tables 1 and 2.

Table 2: Economic Parameters

ρ_y	effective discount rate of young	4.0% per year	$\frac{0.04}{365}$
ρ_o	effective discount rate of old	10% per year	$\frac{0.10}{365}$
$\bar{u}$	value of life	11.4× consumption p.c.	11.24
$\widehat{u}^f$	disutility of fever	lose 30% of baseline utility	-3.37
$\widehat{u}^e$	disutility of emergency care	lose 100% of baseline utility	-11.24
μ^b	size of basic sector	45.4%	0.454
μ^y	share of young	85%	0.85
τ	transfer cost	$0.38 burden of excess taxation	3.51
γ_0	initial share mitigated	50%	0.5
γ_1	speed of mitigation		−0.5
γ_2	time mitigation begins	100 days	100
Θ	hospital capacity	100,000 beds	0.00042
η	bed cost	$7,500	50
λ_o	impact of overuse on mortality	20% higher mortality at peak	6.30

---

[22]These numbers are within the range of expert estimates from the COVID-19 survey compiled by McAndrew (2020) at the University of Massachusetts.

Table 3: Millions of People in Each Health State

	S	A	F	E	R	D × 1000
03/21/20	323.71	4.17	0.84	0.01	1.27	0.34
04/12/20	311.31	2.95	2.72	0.12	12.88	22.1

4   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 in the next section to optimal mitigation.

4.1   Benchmark Results

In Figures 1 to 3, we display the population health dynamics from March 21 to the end of 2020. The red dashed lines represent our baseline scenario with 50% economic mitigation ($m_t = 0.5$) for 100 days and social distancing as described in Section 3.3. The blue solid line is an alternative that shares the same time path of parameters and policies before April 12 – including 50% mitigation between March 21 and April 12 – but in which economic mitigation is set to zero from April 12 onward.
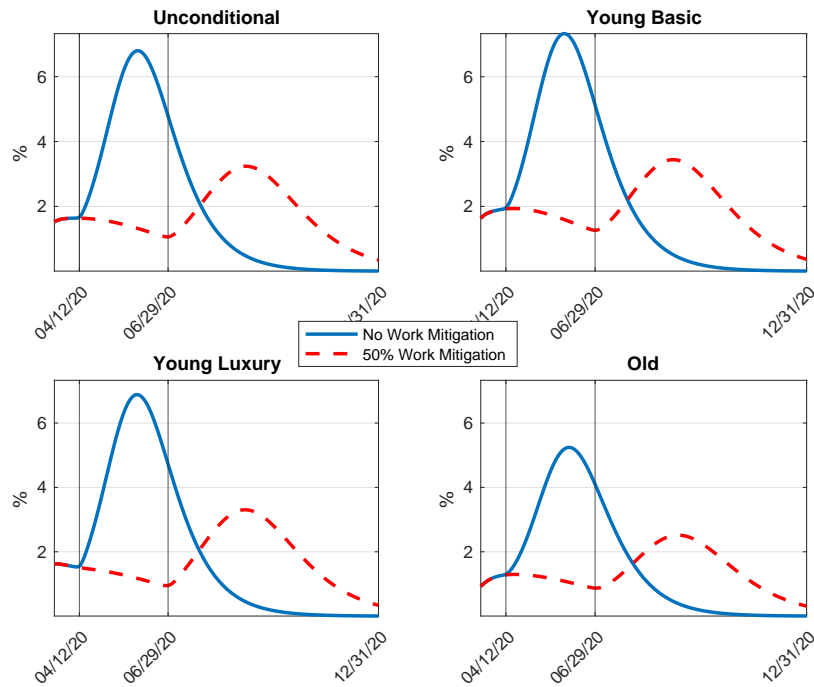

Figure 1: Share of Each Group Infected (Asymptomatic + Fever + Hospitalized).

We start with the currently infected (asymptomatic plus those with fever and those in hospital) in Figure 1. Under our baseline policy, the red dashed line indicates that on April 12, we are already close to a local peak in active infections. In contrast, if economic mitigation were to cease being enforced starting on April 12, the share of the population actively infected would more than triple, reaching almost 7 percent of the population at the end of May. In the scenario when the end of the shutdown is delayed until mid year, the end of mitigation leads to a second wave of infections in the fall, but peak infection rates are much lower than under the scenario when economic mitigation ends now (April 12).

Turning to the heterogeneity across the population, note that absent economic mitigation, basic workers are infected at a slightly higher rate than luxury workers, reflecting the fact that hospitals are in the basic sector. The old – who do not face exposure at work – experience a lower rate of infection than either of the young types. Economic mitigation reduces infection rates for all three types. While mitigation has a larger direct health benefit for luxury workers – they are the ones who stay home from work – all three groups benefit from economic mitigation to a surprisingly similar extent. This is because lower virus spread at work means fewer infected people outside of work and thus fewer new infections at home and in stores and hospitals.

While a smaller share of the old develops mild symptoms, reflecting a lower initial infection rate, a much larger share of the old population ends up being severely sick and hospitalized (Figure 2). This is true under both mitigation scenarios, but the effect is especially pronounced if economic mitigation is abolished early: infections sky-rocket first in the workplace and then at home and during shopping trips, translating into more infections among the old. Recall that conditional on becoming infected, the old are over six times as likely as the young to eventually require hospitalization.[23]

The red horizontal line in the upper left panel of Figure 2 plots hospital capacity, Θ, which we assume to be fixed in the short run. This plot shows another dramatic difference between the two mitigation scenarios. Under the benchmark scenario with 50% economic mitigation until the end of June, the demand for hospital care does not exceed capacity until the fall. In contrast, when economic mitigation policies are (counterfactually) suspended on April 12, capacity is drastically exceeded in May, June and July.

Figure 3 shows daily deaths from COVID-19. Under the baseline policy, with 50% mitigation until mid-year, deaths remain below 2,000 per day until the fall, when the end of mitigation leads to a second wave which peaks at 4,000 deaths per day. When economic mitigation instead ends at Easter, the daily death toll rises dramatically, reaching 11,000 at the peak. The breakdown across population groups indicates that the virus is predicted to kill more older individuals than younger ones, even though the old account for only 15% of the population.

A useful test of the model is to compare model-predicted mortality by age to the data that

---

[23]Figures 12 and 13 in the Appendix plot the incidence of asymptomatic infections and of infections with fever.
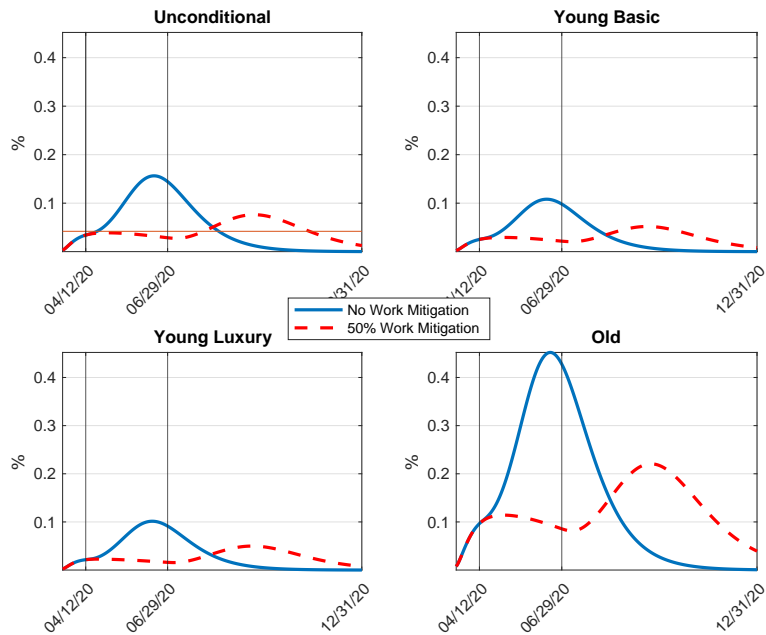
Figure 2: Share of Each Group Hospitalized.

is available to date. In the model, the old account for 73.5% of cumulative deaths up to April 12. By comparison, of the 6,839 deaths reported in New York City as of April 14, 72.3% were associated with individuals over 65.[24] Thus, the age variation in infection and disease progression probabilities in our model is consistent with the observed age distribution of mortality.

In Figure 4 we display the population health dynamics over the next 18 months, starting on April 12. The left panel plots, against time, the share of the population that has not yet been infected (i.e., the susceptible group). The right panel displays the cumulative share of the population that has died from the virus,

Absent economic mitigation, the virus spreads rapidly, and after about six months, 55.4% of the U.S. population has been infected with the virus: the blue line with the never-infected share of the population rapidly drops to 44.6%. In contrast, under our projection for the current economic mitigation plan, the never-infected share declines more slowly, and a larger share of the population is never touched by the virus (51.5% rather than 44.6%). That is, 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

---

[24]Data from New York City Health Department as reported by Worldometers https://www.worldometers.info/coronavirus/coronavirus-age-sex-demographics/
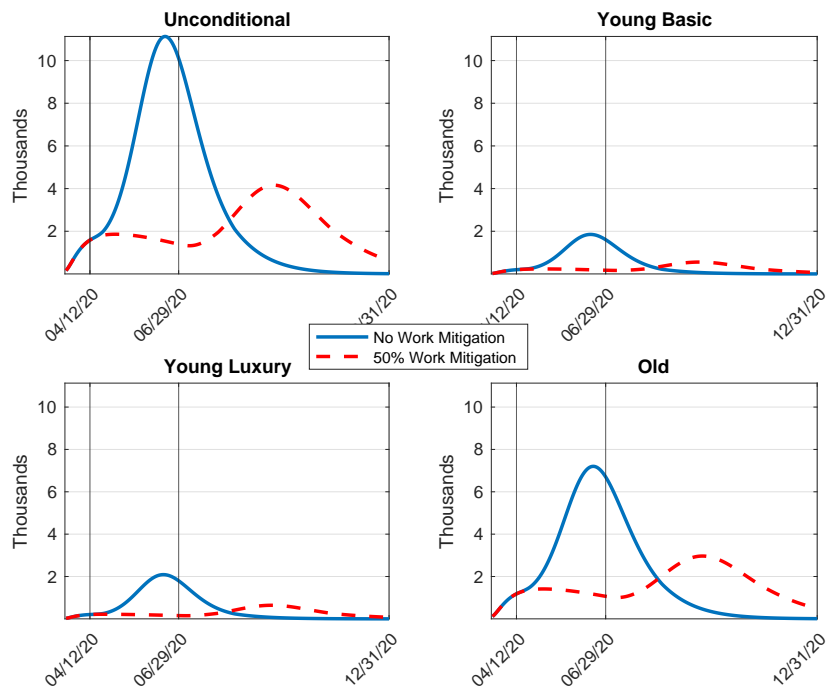
Figure 3: Daily Deaths from COVID-19

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 the economy converges to a steady state in which a larger share of individuals has never been infected, relative to the scenario in which the economy open up at Easter.

The right panel translates infections into mortality associated with the virus. In the absence of economic mitigation, the death toll of the virus rises rapidly, and by the end of the outbreak 0.26% of the U.S. population is predicted to have lost their lives, which amounts to $858,000$ people. Under the current benchmark economic mitigation policy, that number falls to 0.19% ($627,000$ individuals). The difference in lives lost ($231,000$) comes from two sources. First, with economic mitigation in place, there is less hospital overload and excess associated mortality. Second, with mitigation, a smaller cumulative total number of infections means that fewer people ever risk adverse health outcomes and death. Of the $858,000$ total death toll absent any economic mitigation from April 12 onward, $191,000$ deaths are due to hospital capacity being exceeded. Under the baseline 50-percent-for-100-days mitigation policy, only $32,700$ out of $627,000$ deaths reflect hospital overload. Thus, $158,300$ of the extra $231,000$ lives lost when
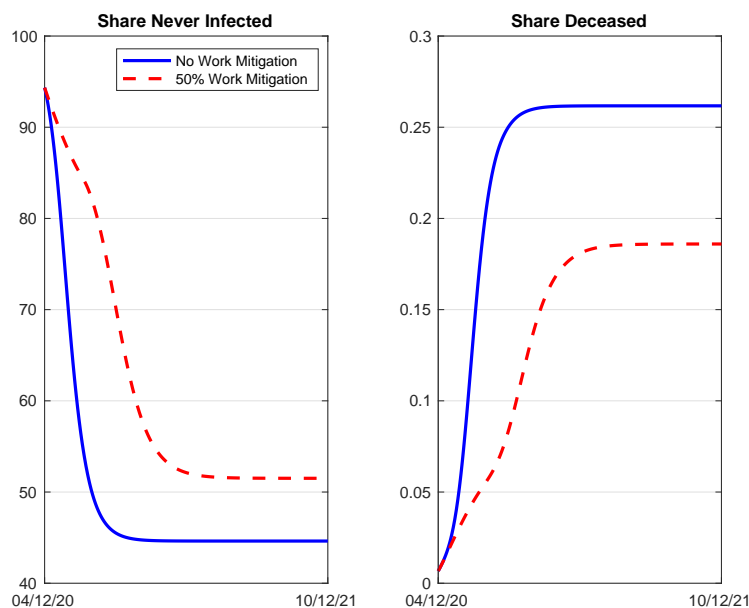
25

Figure 4: Left Panel: Share of the Initial Population Deceased. Right Panel: Share of Population Never Infected (Susceptible).

the shutdown ends at Easter reflect a severely over-stretched hospital system.[25]

Figure 5 plots the dynamics of consumption for workers, and non-workers through the course of the pandemic. Recall that in this economy, all workers independent of sector, enjoy the same consumption level, 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 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 freely redistribute. In that case, the planner equates consumption between workers and non-workers at each date.[26]

Comparing across columns, the left two panels display the evolution of consumption when economic mitigation ends on April 12, and the right two panels maintain 50% mitigation until the end of June. In the first case, the economy immediately recovers as all healthy workers who were affected by the shutdown in the luxury sector return to work, increasing output,

---

[25]Table ?? in the Appendix reports the number of people forecast to be in each health state at various different dates, under the baseline 50-percent-for 100-days mitigation policy.

[26]Recall that the evolution of the population health distribution is independent of the cost of transfers.
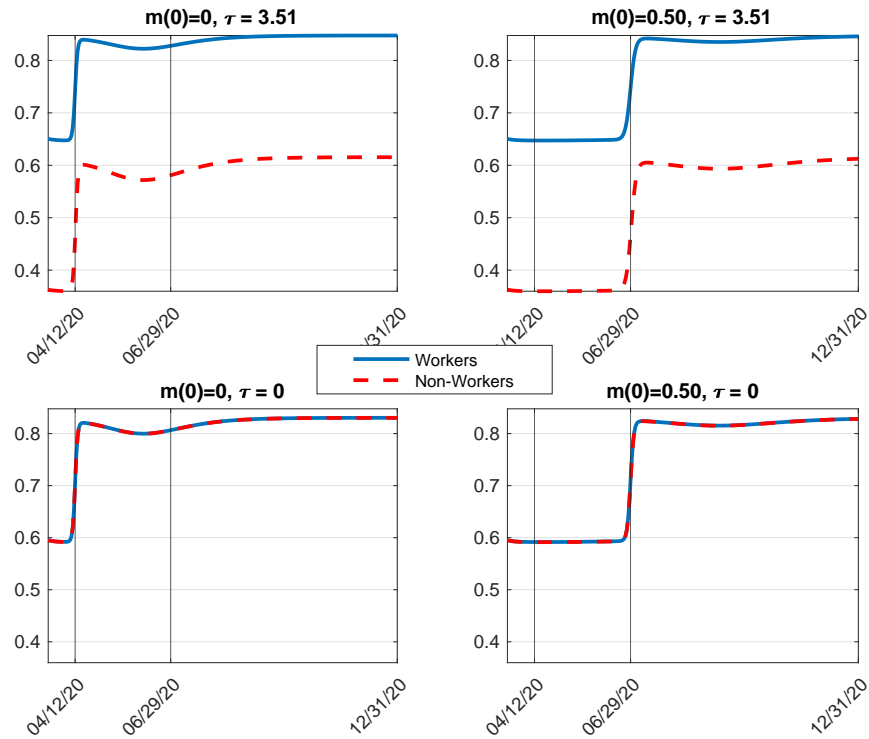
Figure 5: Consumption Paths. Top Two Panels, $\tau = 3.51$. Bottom Two Panels, $\tau = 0$. Left Two Panels, $m = 0$. Right Two Panels, $m = 0.5$ for 100 Days, Then $m = 0$.

income, and thus aggregate consumption in the economy by about 27.5%.[27] The right two panels show that in terms of output and thus consumption, a later end to the shutdown simply (and somewhat mechanically) postpones the economic recovery by 2.5 months. Note from the upper right panel of Figure 5 that the cost of economic mitigation is borne 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 and increases it again as the economy recovers.

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

---

[27]Note that we assume that infected people with symptoms stay home rather than go to work, and since the share of infected individuals is endogenously evolving over time, the increase is not exactly equal to the 27.5% decline in output when economic mitigation was introduced in the first place.

Table 4: Welfare Gains (+) or Losses (-): Mitigation from 4/12/20–6/29/20

Mitigated Share	75%		50%		25%	
Transfer Cost τ	3.51	0	3.51	0	3.51	0
Young Basic	0.06%	-0.04%	0.24%	0.18%	0.33%	0.30%
Young Luxury	-0.37%	-0.05%	-0.01%	0.16%	0.23%	0.29%
Old	1.44%	2.00%	2.17%	2.64%	2.60%	2.93%

of redistribution. In particular, we consider three mitigation levels: $m = 0.5$ (our baseline used to construct the previous plots), $m = 0.75$, and $m = 0.25$. The welfare calculation asks, What percent of consumption would a person be willing to pay every day for the rest of her life to move from the economy where work mitigation ends on April 12, 2020 to one where work mitigation changes to $m = 0.75$ or $m = 0.25$ (or remains at $m = 0.5$) through June 29, 2020? For this calculation, we use April 12 as the starting date, and assume m is fixed at the values considered level until June 29, after which date $m = 0$ in each case. We report results for our baseline value for τ (3.51) and for a case in which redistribution is costless ($\tau = 0$).

The first clear message from Table 4 is that economic mitigation offers significant welfare gains for the old but has much more modest welfare effects on the young. For example, in our baseline case ($m = 0.5$ and $\tau = 3.51$), the old gain 2.17% of consumption, while the young basic workers gain only 0.24% from the shutdown, and young luxury workers are marginally worse off. 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, and strong economic mitigation policies reduce infections in the workplace, which in turn lowers the risk that the old meet infected individuals at home or while shopping.

The second key message is that the cost of redistribution matters. In particular, when redistribution is costless, young luxury workers and young basic workers perceive essentially identical welfare effects from mitigation.[28] However, when redistribution is costly, young luxury workers fare notably worse than young basic workers because they risk larger expected 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.

---

[28]On the one hand, mitigation offers more direct protection to luxury workers, because they are the ones to stay home. On the other hand, mitigation reduces hospitalizations, which reduces transmission to basic sector hospital workers. These two effects essentially offset.

4.2 Optimal Policy

The mitigation policies we have compared thus far were not chosen optimally. We now turn to exploring the optimal time path for economic mitigation and the associated statically optimal degree of redistribution, given that path. To start, we optimize over the three parameters in our parametric process for m_t. That is, we choose γ_0, γ_1, and γ_2 in eq. 30 to maximize social welfare as defined in Section 3.2. The choice of these parameters lets the government control, respectively, the initial size of economic mitigation, when it ends, and how quickly it is phased out. We have also computed the optimal non-parametric path for m_t. This more flexible specification turns out to deliver negligible welfare gains relative to our parametric function (see section 4.3). We therefore focus on the parametric case, given its advantages in terms of communication and interpretation.

Figure 6 describes the preferred policies within this class. The left panel describes optimal policies under our baseline cost for redistribution, with $\tau = 3.51$. The blue line is the policy chosen by a utilitarian planner, who weighs the expected utility of each type in proportion to its date 0 population shares. The other lines describe the policies preferred by each of the three different types (young workers in the basic sector, young workers in the luxury sector, and old individuals, respectively). The right panel corresponds to a case in which redistribution to soften the economic effects of mitigation is costless ($\tau = 0$).

There are clearly large differences across individual types in terms of what fraction of the economy they would like to see shut down and for how long. As a point of comparison, recall that up until April 12, the level of mitigation is set at 50% of the luxury sector. We first focus on the benchmark calibration with costly transfers (the left panel).

The old (15% of the population) would like to see 30% of the luxury sector shut down, and for a partial shutdown to remain in place through the end of the year. In contrast, young luxury workers (close to 50% of the entire population) would prefer a much lower level of mitigation and for that mitigation to end much earlier. Basic sector workers have a policy preference roughly in the middle of these two extremes, and a utilitarian government adopts a similar policy. Thus, a utilitarian government closes about 25% of the luxury sector until around mid July, before gradually opening up over the following couple of months. Note that this policy implies a notably lower level of economic mitigation than the one currently in place, but indicates that mitigation should remain in place for longer than our baseline 100 day assumption.

When redistribution is costless (right panel of Figure 6), policy preferences remain qualitatively similar but quantitatively change significantly. First, young workers in both sectors now agree on the preferred mitigation policy, which is because they face identical consumption consequences, and benefit essentially equally on the health front.[29] Second, the old now prefer even

---

[29] Although again mitigation benefits the two groups via different channels: reduced hospital infection for basic workers, and reduced workplace infection for luxury workers.
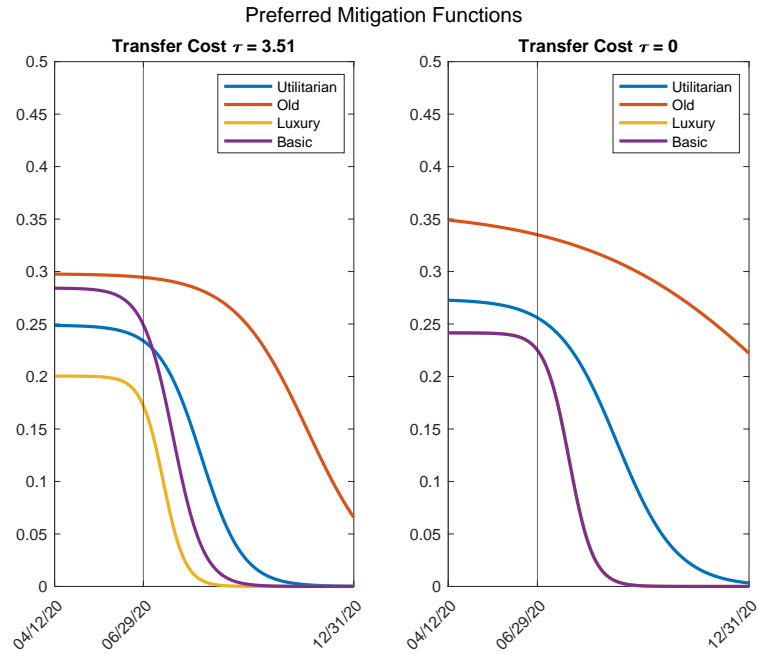
Figure 6: Preferred Intensity and Duration of Shutdowns

Table 5: Welfare Gains (+) or Losses (-) from Preferred Mitigation, $\tau = 3.51$

	Utilitarian	Old	Luxury	Basic
Young Basic	0.37%	0.29%	0.34%	0.36%
Young Luxury	0.21%	-0.05%	0.25%	0.22%
Old	3.60%	4.15%	2.89%	3.37%

more mitigation, because they do not have to worry about a reduction in relative consumption during a shutdown. The utilitarian policy is more aligned with the preferences of young workers, simply because they constitute the lion's share of the population. Importantly, the preferred utilitarian mitigation policy is more aggressive when redistribution is (counter-factually) cost-less, both in terms of level as well as in terms of a longer and more gradual phasing-out. Thus, heterogeneity coupled with costly redistribution reduces the pay-off to using shutdowns as a tool to fight the pandemic.

The next two tables (Tables 5 and 6) describe expected welfare gains, relative to a no-economic-mitigation baseline, under each of the policies described in Figure 6. In each case, the starting date for these welfare evaluations is April 12. The columns of each table identify the policy in place. The rows report expected welfare for each type.

Table 6: Welfare Gains (+) or Losses (-) from Preferred Mitigation, $\tau = 0$

	Utilitarian	Old	Luxury	Basic
Young Basic	0.30%	-0.05%	0.32%	0.32%
Young Luxury	0.29%	-0.06%	0.32%	0.32%
Old	4.49%	5.30%	3.68%	3.68%

Consistent with the results in the previous section, the old experience large welfare gains from any of these policies. Irrespective of the cost of redistribution, the welfare gains or losses for the young are much smaller. Second, and again in line with the previous section, the welfare gains for young luxury workers are always smaller than for young basic workers when redistribution is costly, but nearly identical when redistribution is costless. Third, the policies that are welfare-maximizing for the old can be welfare-reducing (relative to no mitigation) for the young, especially those in the luxury sector.

Finally, Figure 7 compares key model predictions under three policies: (1) 50 percent mitigation for 100 days from March 21 (red), (2) ending mitigation on April 12 (dark blue), and (3) the optimal policy chosen by a utilitarian planner (cyan). The top left panel highlights the key optimal policy finding: the shutdown should be relatively modest (initially 25 percent of the luxury sector or 13.7 percent of the economy as a whole) but it should stay in place for a long time. This translates into a modest but prolonged contraction in output and consumption (top right panel). This policy delivers a mild but lengthy increase in infections, hospitalizations, and deaths during the spring. But crucially there is no second wave of infections in the fall. A final and very important benefit of the optimal policy is evident in the bottom right panel: it translates into a larger share of the population never being touched by the virus. In particular, the share never infected by the end of 2021 under the optimal policy is 58.1 percent under the optimal policy, compared to 51.5 under the harsh 100 day shutdown, and 44.6 under the counterfactual in which the shutdown ends at Easter.

These different dynamics for infections and hospitalizations translate into large differences in cumulative deaths under the different policies. The total number of deaths under the utilitarian optimal policy is $533,400$, compared to $621,900$ under the policy when the shutdown ends at the end of June, and $868,800$ when it ends at Easter.

4.3 Small Gains From More Complex Policies

We now solve for the optimal path of mitigation without restricting the path for mitigation in any way, i.e., we solve for the fully optimal non-parametric path. This amounts to numerically solving the optimal control problem outlined in Appendix 1. We then address two questions. First, how different are the fully optimal paths of mitigation and induced variables relative to
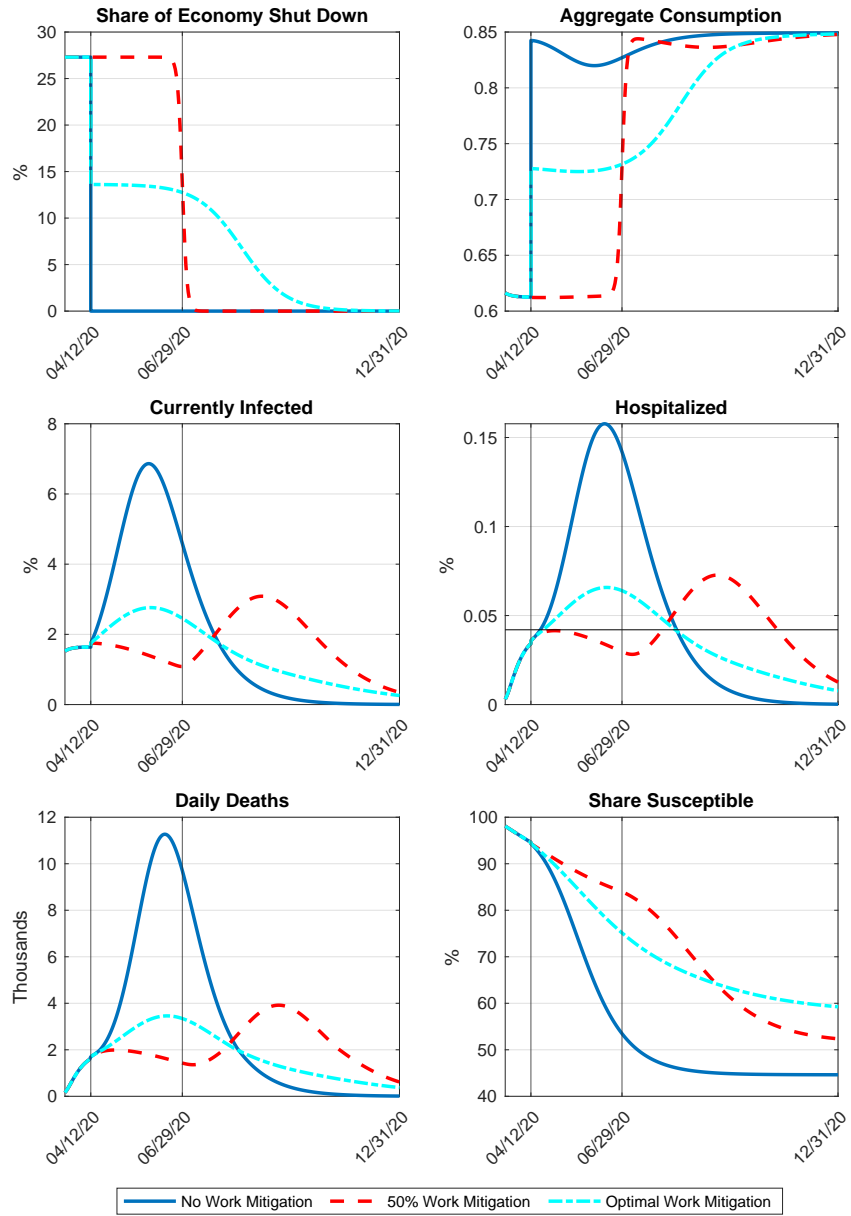
Figure 7: Key Outcomes under Alternative Mitigation Policies

the best-in-parametric-class policy we have focused on to date? Second, how much better can the planner do, in welfare terms, when she can choose any possible path for mitigation instead of being restricted to our baseline parametric function of time?

Figure 8 illustrates the optimal paths of key variables when the planner chooses the path for mitigation optimally with no parametric restrictions, and offers a comparison to the corresponding paths under our optimal parametric mitigation policy. The top left panel shows the amount of economic activity that is shut down under each policy. The optimal non-parametric policy tracks the parametric policy closely, though the extra flexibility allows for a hump shape in mitigation that tracks infections more closely than the simple policy.

The more flexible policy implies slightly more volatility in consumption, which reduces welfare, while slightly flattening the curves for infections and hospitalizations, which increases welfare. Overall, Table 7 indicates tiny marginal welfare gains from the optimal non-parametric policy relative to the optimal parametric one: flexibility delivers a 0.01% increase in the utilitarian planner's payoff.

We have also experimented with computing optimal non-parametric policies starting on February 15, as opposed to April 12. With an earlier start date for policy, the additional welfare gains from allowing a more flexible path for mitigation are much larger. In particular, given a February 15 start date, it is optimal to impose relatively modest mitigation for a while, before ramping up (see Figure 14 in the Appendix). This sort of policy is precluded by our parametric functional form, which therefore delivers welfare inferior outcomes.

Table 7: Welfare Gains (+) or Losses (-): Non-Parametric vs. Parametric Policies

Policy Form	Utilitarian		Old		Luxury		Basic	
	Non-Par	Par	Non-Par	Par	Non-Par	Par	Non-Par	Par
Young Basic	0.37%	0.36%	0.30%	0.29%	0.34%	0.34%	0.37%	0.37%
Young Luxury	0.22%	0.21%	-0.03%	-0.05%	0.25%	0.25%	0.22%	0.22%
Old	3.55%	3.60%	4.15%	4.15%	2.79%	2.89%	3.38%	3.37%

We now briefly discuss a few factors that shape our results on optimal mitigation. First, the overall level of the welfare numbers is sensitive to several choices. A key one is the value of a statistical life: a lower value would make life-saving economic mitigation trivially less attractive. Second, if we assumed lower recovery rates at different stages of an infection, or a higher mortality rate at the hospital stage, agents would perceive a greater risk of death and 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
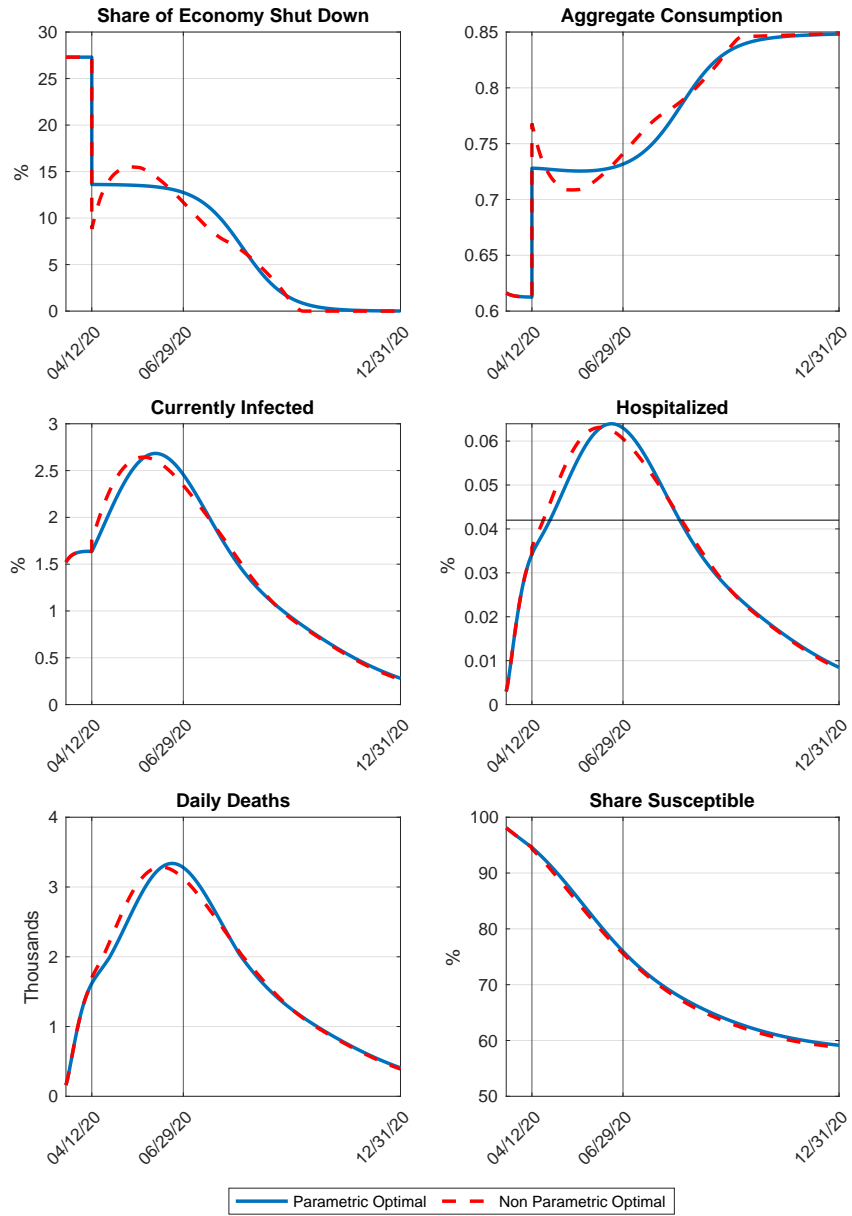
Figure 8: Key Outcomes, Parametric Vs Non-Parametric Utilitarian Policies

compensates strongly for the utility cost of reduced consumption. Fourth, if business failures caused by the shutdown lead to the destruction of firm-specific capital, then the shutdown will continue to depress output even after it is lifted, suggesting a softer optimal shutdown. Finally, the attractiveness of shutdowns clearly depends on the share of virus transmission that occurs through different forms of economic activity: the larger that share is, the more powerful shutdowns are as a tool to slow transmission.

5 Extensions

We now consider how three different public health innovations would affect optimal mitigation and redistribution policies. First, we consider a case in which a successful vaccine is expected to arrive at the end of 2020. This implies large benefits to avoiding infections until the vaccine arrives, which translates into much more extensive mitigation than in our baseline simulation in which there is no prospect of a vaccine. Second, we consider the possibility that rapid testing for antibodies becomes available, so that all recovered workers can be identified and exempted from mitigation. In this case, we find slightly more desire for mitigation, since it can be better targeted. Overall, however, antibody testing has disappointingly little impact on the trajectory of the pandemic or optimal policy. Third, we consider an intervention in which the young do all the shopping for the old, allowing the old to be protected from infection through this channel. This reduces optimal mitigation for a utilitarian planner, with additional welfare benefits for the old.

5.1 Policy Response to a Future Vaccine

Researchers around the world are racing to develop a COVID-19 vaccine. A team at the Jenner Institute at the University of Oxford hopes to have one in production by the end of 2020.[30] How do preferred policies change if people know that a vaccine is in the pipeline? To answer this question we solve for optimal policies when individuals expect a perfectly effective vaccine to end the flow of new infections on December 31, 2020.

Figure 9 illustrates the optimal policy and key outcomes when a vaccine is expected. The main takeaway is that the planner prefers a much more extensive shutdown when she knows a vaccine is coming. The reason mitigation is now more attractive is that the arrival of a vaccine stops new infections cold, and thus mitigation can dramatically reduce the total cumulative number of infections and deaths from COVID-19. Absent a vaccine, in contrast, mitigation is less attractive, because most of the infections that mitigation prevents in the short run are simply postponed further into the future. Effectively, without a vaccine the epidemic will only die down after some form of herd immunity has been achieved. Mitigation efforts can only delay deaths and avoid an overload of the health system. In contrast, if a vaccine is expected at some

---

[30]https://www.bbc.com/news/uk-53469839 Note that most experts do not expect a vaccine to be widely available until 2021.
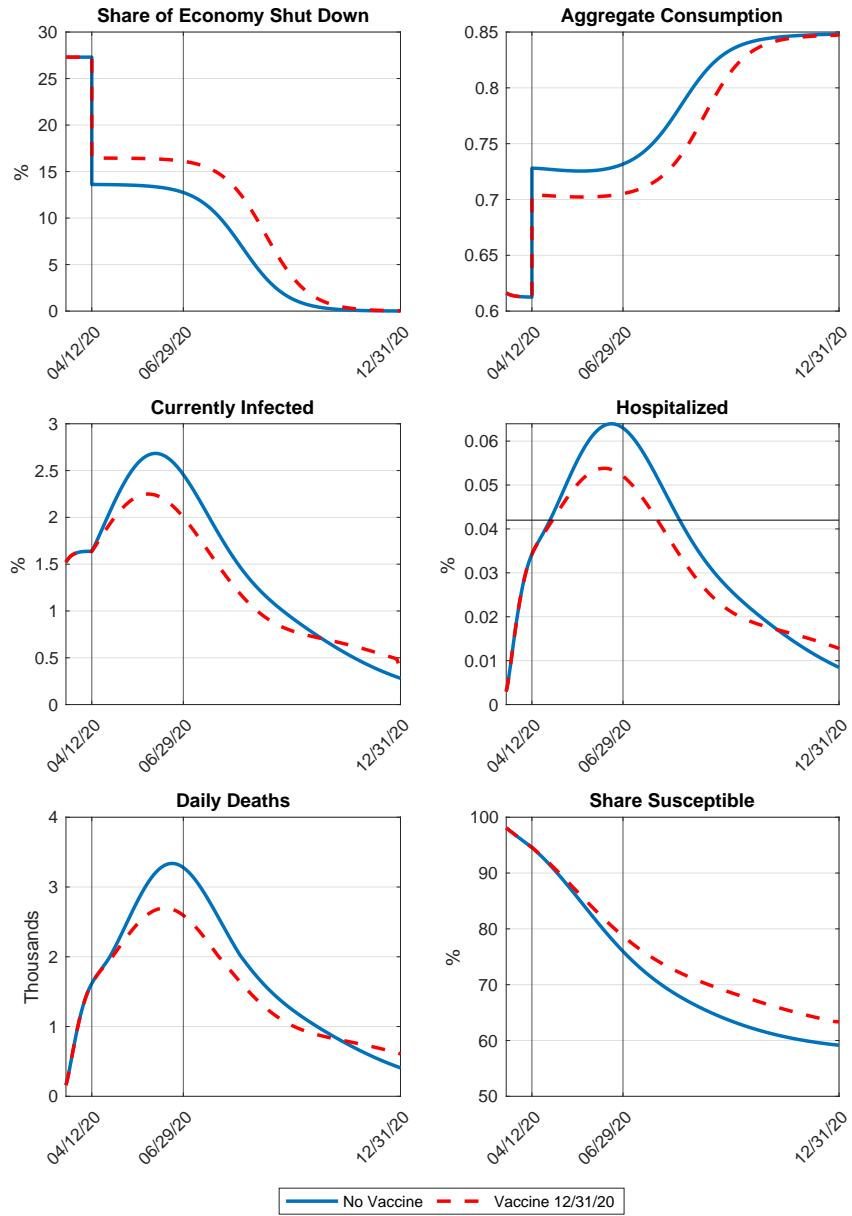
Figure 9: Key Outcomes, Baseline vs. Vaccine December 31 2020

date, postponing potential infections until after that date massively increases the share of the population that never contracts the disease (see the bottom right panel).

The ability to mitigate until a vaccine arrives and therefore avoid many serious illnesses and deaths implies large welfare gains from mitigation, especially for the old (see column 3 of Table 8). The welfare gains for the old nearly double in the scenario in which a vaccine will arrive, relative to the no-vaccine baseline. Basic sector workers also gain more from mitigation relative to the baseline no-vaccine case. In contrast, luxury sector workers view the effects of mitigation similarly whether or not they expect a vaccine since their overall risk of getting infected, severely sick and die from the disease is relatively low.

5.2   Policy Response to Antibody Testing

The mitigation of luxury workers studied thus far has been uniform across susceptible, asymptomatic, and recovered people. Given that we maintain the assumption that recovered people are no longer infectious, nor can be reinfected, it would make sense to allow them back to work. The challenge is that identifying who is recovered is not easy, for two reasons. First, people in the model (and in the world) can recover from the asymptomatic state, and this group might not be aware of their own recovered status. Second, even if people are privately fully aware of whether they are in either (i) the susceptible or asymptomatic state, or (ii) are recovered, they might individually prefer to work rather than to be mitigated, given that workers enjoy higher consumption.[31] Thus, identifying health status would remain a challenge for the planner.

However, in recent weeks, antibody testing has become much more widely available. We now consider a scenario in which the planner tests all young workers at high frequency, and offers immunity passports to all recovered individuals, exempting them from mitigation. Figure 10 plots key outcomes under the utilitarian-optimal parametric policy in the economy with antibody testing.

The planner now chooses slightly more extensive and notably longer-lasting mitigation (top left panel). Changes in outcomes are relatively small. Column (4) of Table 8 shows that each group gains more from the utilitarian optimal policy if antibody testing is available, especially the old who gain 0.31% in consumption equivalent terms. Note, however, that these welfare gains are much smaller than those associated with the arrival of a vaccine.

This finding of relatively modest welfare gains might seem surprising and disappointing, given that we have assumed maximally effective antibody testing. The intuition for why antibody testing is not especially useful is as follows. In the early days of the pandemic, there are very few recovered individuals, so being able to identify them is not very helpful. Toward the end of the pandemic, exempting a large number recovered from lockdowns is more useful, because

---

[31]We have verified that under the consumption allocation characterized in Section 3.2 the young in our model do prefer working to being mitigated.
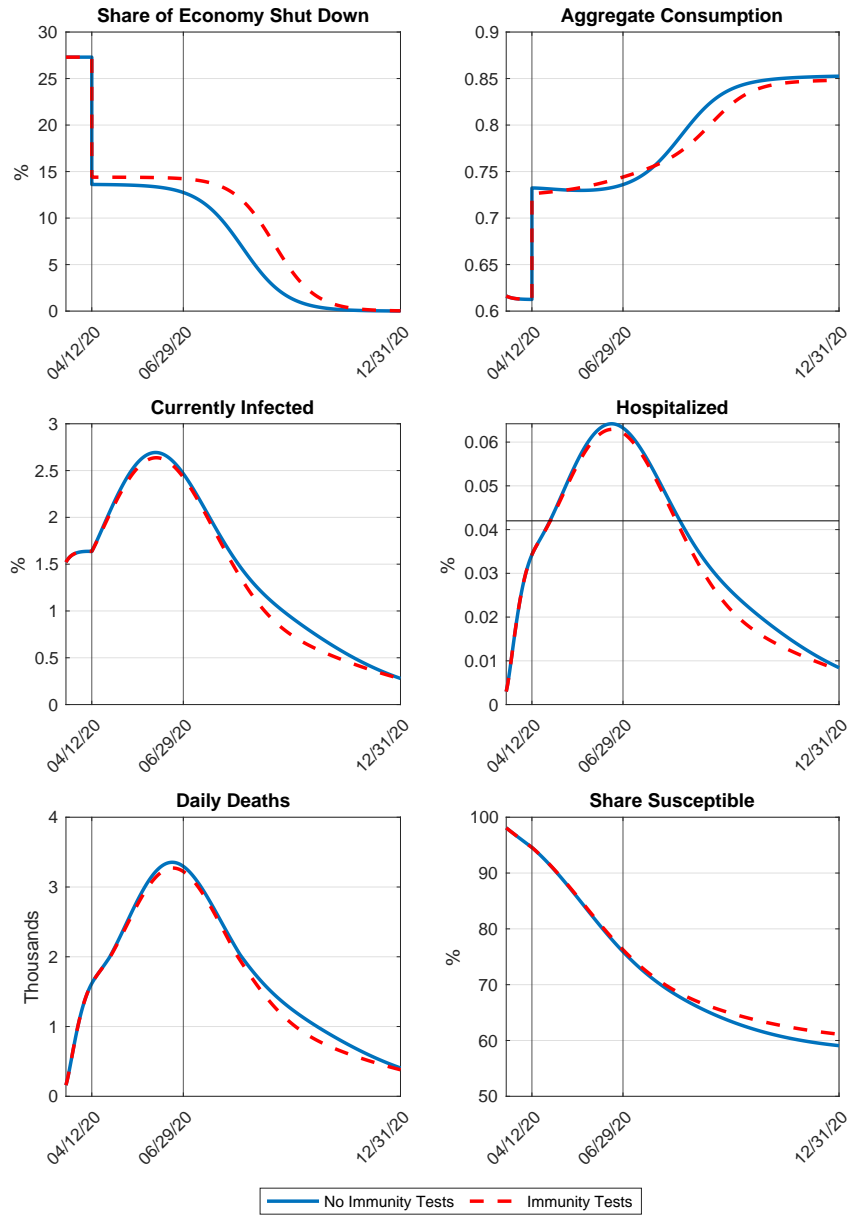
Figure 10: Key Outcomes, Baseline vs. Antibody Testing

it implies that a significantly larger share of those mitigated are asymptomatic (or susceptible) rather than recovered. For this reason mitigation is optimally prolonged. However, in this phase the disease is dying out anyway, so the welfare gains from further hastening the end of the pandemic, while real, are fairly modest.

5.3 Policy Response to Elderly Isolation

One widely-proposed policy has been to isolate older people at home. In our model, the old do not work, so the only way to lower their risk of infection outside the home is to reduce their market consumption activity. We therefore consider a simple intervention to protect the old from infection while shopping, which is to ask the young to shop on behalf of the old.[32]

To model this, we remove the consumption channel as a possible source of infection for the old, and simultaneously increase the importance of the consumption channel for the young, in such a way that the total number of infections through consumption would remain unchanged, given the same distribution across health states for young and old.[33] Importantly, we assume that this policy intervention is costless.

Obviously this intervention directly benefits the old. But, once the old do not shop, the planner wants to change the path for mitigation, because mitigation no longer directly protects the old from infection. In particular, the planner now optimally reduces mitigation (see the top left panel of Figure 11). This benefits young workers, and especially those in the luxury sector.

Figure 11 illustrates that reduced mitigation translates into more infections. But because these infections are more heavily tilted toward the young, deaths are reduced. Fewer deaths plus higher average consumption is an attractive package. Column (5) of Table 8 documents welfare gains for all groups. In this instance we report welfare gains from a joint policy of isolation of the old plus optimal mitigation relative to a baseline of no mitigation and no isolation. The table indicates similar welfare gains for the young and even larger welfare gains for the old, relative to the baseline case (column 1) in which the isolation instrument is not used.

Table 8: Welfare Gains (+) or Losses (-): Utilitarian Planner in Extensions

	Baseline	Non-Parametric	Vaccine	Antibody Tests	Elderly Shut In
Young Basic	0.36%	0.37%	0.39%	0.38%	0.34%
Young Luxury	0.21%	0.22%	0.16%	0.23%	0.20%
Old	3.60%	3.55%	4.52%	3.91%	4.80%

---

[32] Alternatively, this policy can be interpreted as introducing special extended shopping hours for the elderly where infection risk in minimized, in turn reducing shopping hours for the young, resulting in more crowded and infectious stores for them.

[33] In particular, the consumption term in eq. (4) becomes $\beta_c(m)y(m)x^{ya}\left(\frac{x^y+x^o}{x^y}\right)x^{ybs}\left(\frac{x^y+x^o}{x^y}\right)$. The adjustment for eq. (5) is similar.
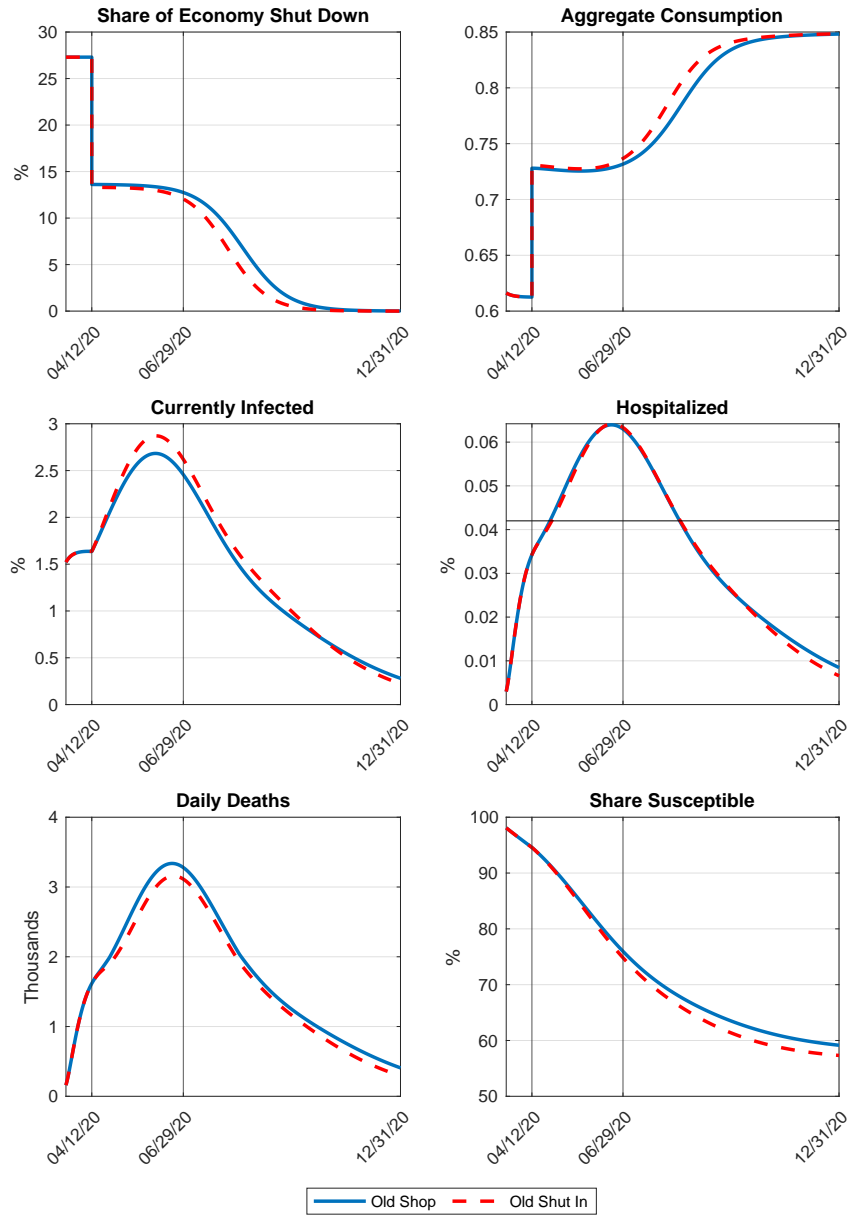
Figure 11: Key Outcomes, Baseline vs. Elderly Shut In

6 Conclusion

A key challenge in designing an optimal policy response to a pandemic like COVID-19 is that mitigation efforts offer large potential benefits to some groups (the old) while imposing large costs on others (the young, and especially those whose jobs are shuttered). The fact that the gains and losses from mitigation are unequally distributed makes fighting pandemics difficult. Mitigation efforts are likely to be more popular when the costs of shutdowns can be distributed more evenly across the population via redistribution. However, redistribution is costly, so mitigation efforts may not be Pareto-improving in practice.

Within a quantitative heterogeneous agent model, we have computed an optimal joint shutdown and redistribution policy. Relative to predictions from models without heterogeneity or with costless redistribution, the utilitarian optimal shutdown is relatively mild, at the cost of more adverse health outcomes. Our baseline calibration suggests that the shutdown in place on April 12 was too extensive, but that a utilitarian planner would keep a partial shutdown in place into the fall.

To the extent that countries differ in terms of the efficiency of tax and transfer systems, our framework predicts more aggressive shutdown efforts in countries with more efficient redistribution systems. The model predicts extensive disagreement about how extensive shutdowns should be, and how long they should last. Thus, the actual policy implemented will depend on the political power of different constituencies. Modeling the joint politico-economic determination of shutdown and redistribution policies is an important avenue for future research.

References

ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. WHINSTON (2020): "Optimal Targeted Lockdowns in a Multi-Group SIR Model," Working Paper, MIT.

ARGENTE, D., F. ALVAREZ, AND F. LIPPI (2020): "A Simple Planning Problem for COVID-19 Lockdown and Smart Tracing," Working Paper, University of Chicago.

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.

BICK, A. AND A. BLANDIN (2020): "Real Time Labor Market Estimates During the 2020 Coronavirus Outbreak," Unpublished Manuscript, Arizona State University.

BROTHERHOOD, L., P. KIRCHER, C. SANTOS, AND M. TERTILT (2020): "An Economic Model of the Covid-19 Pandemic with Young and Old Agents: Behavior, Testing and Policies," Working Paper, University of Mannheim.

CAJNER, T., L. D. CRANE, R. A. DECKER, A. H.-P. JOHN GRIGSBY, E. HURST, C. KURZ, AND A. YILDIRMAZ (2020): "The U.S. Labor Market During the Beginning of the Pandemic Recession," Working Paper, University of Chicago.

CHARI, V., R. KIRPALANI, AND C. PHELAN (2020): "The Hammer and the Scalpel: On the Economics of Indiscriminate versus Targeted Isolation Policies during Pandemics," Working Paper, University of Minnesota.

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.

FARBOODI, M., G. JAROSCH, AND R. SHIMER (2020): "Internal and External Effects of Social Distancing in a Pandemic," University of Chicago, Working Paper.

FERGUSON, N. M. ET AL. (2020): "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand," Imperial College COVID-19 Response Team.

FERNANDEZ-VILLAVERDE, J. AND C. JONES (2020): "Estimating and Simulating a SIRD-Model of COVID-19 for Many Countries, States, and Cities," Working Paper, Stanford University.

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.

GUERRIERI, V., G. LORENZONI, L. STRAUB, AND I. WERNING (2020): "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" Northwestern University, Working Paper.

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," Centre for Economic Policy Research DP 14607.

MCANDREW, T. (2020): "COVID19-Survey 8-2020-04-08," Unpublished Paper, University of Massachusetts Amherst.

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.

SAEZ, E., J. SLEMROD, AND S. H. GIERTZ (2012): "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature*, 50, 3–50.

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

TSAI, T. C., B. H. JACOBSON, AND A. K. JHAJACK (2020): "American Hospital Capacity And Projected Need for COVID-19 Patient Care," Health Affairs Blog, March 17, 2020. DOI: 10.1377/hblog20200317.457910.

XU, J., G. WATSON, D. XIONG, M. A. SUCHARD, AND C. M. RAMIREZ (2020): "COVID-19 Infection and Mortality across Workplaces and Intervention Effects," Mimeo, UCLA.

A Details of the SAFER Model

This section summarizes the details of the remainder of the epidemiological block. It is fairly mechanical and simply describes the transition of individuals though the health states (asymptomatic, fever-suffering, hospitalized, and recovered) once they have been infected. Equations (31) to (33) describe 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 developing a fever occurs at rate σ^{yaf} (σ^{oaf}) for the young (old), and exit to the recovered state occurs at rate σ^{yar} (σ^{oar}) for the young (old). Note that someone who recovers at this stage will never know that she contracted the virus.

For individuals suffering from a fever, equations (34) to (36) show that for the young there is entry from the asymptomatic state and exit to the hospitalized state at rate σ^{yfe}, and to the recovered state at rate σ^{yfr}, with analogous expressions for the old. Equations (37) to (39) describe the movements of those in emergency care, showing entry from those with a fever and exits to death and recovery. The death rate is $\sigma^{yed} + \varphi$, while the recovery rate is $\sigma^{yer} - \varphi$, where φ, described in equation 11 in the main text, is a term related to hospital overuse. Equations (40) to (42) display 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 they play no further role in the model.

To summarize, the dynamic system of health transitions from the asymptomatic to the recovered (and death) state is then given by:

$$\dot{x}^{yba} = -\dot{x}^{ybs} - \left(\sigma^{yaf} + \sigma^{yar}\right) x^{yba} \tag{31}$$

$$\dot{x}^{y\ell a} = -\dot{x}^{y\ell s} - \left(\sigma^{yaf} + \sigma^{yar}\right) x^{y\ell a} \tag{32}$$

$$\dot{x}^{oa} = -\dot{x}^{os} - \left(\sigma^{oaf} + \sigma^{oar}\right) x^{oa} \tag{33}$$

$$\dot{x}^{ybf} = \sigma^{yaf} x^{yba} - \left(\sigma^{yfe} + \sigma^{yfr}\right) x^{ybf} \tag{34}$$

$$\dot{x}^{y\ell f} = \sigma^{yaf} x^{y\ell a} - \left(\sigma^{yfe} + \sigma^{yfr}\right) x^{y\ell f} \tag{35}$$

$$\dot{x}^{of} = \sigma^{oaf} x^{oa} - \left(\sigma^{ofe} + \sigma^{ofr}\right) x^{of} \tag{36}$$

$$\dot{x}^{ybe} = \sigma^{yfe} x^{ybf} - \left(\sigma^{yed} + \sigma^{yer}\right) x^{ybe} \tag{37}$$

$$\dot{x}^{y\ell e} = \sigma^{yfe} x^{y\ell f} - \left(\sigma^{yed} + \sigma^{yer}\right) x^{y\ell e} \tag{38}$$

$$\dot{x}^{oe} = \sigma^{ofe} x^{of} - \left(\sigma^{oed} + \sigma^{oer}\right) x^{oe} \tag{39}$$

$$\dot{x}^{ybr} = \sigma^{yar} x^{yba} + \sigma^{yfr} x^{ybf} + (\sigma^{yer} - \varphi) x^{ybe} \tag{40}$$

$$\dot{x}^{y\ell r} = \sigma^{yar} x^{y\ell a} + \sigma^{yfr} x^{y\ell f} + (\sigma^{yer} - \varphi) x^{y\ell e} \tag{41}$$

$$\dot{x}^{or} = \sigma^{oar} x^{oa} + \sigma^{ofr} x^{of} + (\sigma^{oer} - \varphi) x^{oe} \tag{42}$$

$$\tag{43}$$

B Flow Utilities for Each Household Type

Given the consumption allocations characterized in the main text expected flow utility for each household type is given by:

$$W^\ell(x, m) \;=\; \frac{(x^{y\ell n} + x^{y\ell e} + x^{y\ell r})}{x^\ell}\left[(1-m)u(c^w) + mu(c^n) + \bar{u}\right]$$

$$+\frac{(x^{y\ell f} + x^{y\ell e})}{x^\ell}\left[u(c^n) + \bar{u} - \hat{u}\right]$$

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

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

C The Unrestricted Optimal Policy Problem

The complete characterization of the optimal policy path is the solution to an optimal control problem. In the main text we already 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 $\mathbf{x} = (x^{i,j})$) evolves according to the vector-valued equation (summarizing Equations (4) to (42) the paper in a compact form):

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

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

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

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

$$W_m(\mathbf{x}, m) \;=\; -\boldsymbol{\mu} \cdot G_m(\mathbf{x}, m) \tag{46}$$

$$\dot{\boldsymbol{\mu}} \;=\; \rho\boldsymbol{\mu} - \left[W_{\mathbf{x}}(\mathbf{x}, m) + \boldsymbol{\mu} \cdot G_{\mathbf{x}}(\mathbf{x}, m)\right] \tag{47}$$

$$\dot{\mathbf{x}} \;=\; G(\mathbf{x}, m) \tag{48}$$

The key trade offs with mitigation efforts m discussed in the main text are encoded in equation (46). A marginal increase in m entails static economic costs of $W_m(\mathbf{x}, m)$ 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 $G_m(\mathbf{x}, m)$. Concretely, as is clear from equations $(4-6)$ 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 $(\mathbf{x}, \boldsymbol{\mu})$ are vectors, so are the entities $G_m(\mathbf{x}, m) = (G_m^{i,j}(\mathbf{x}, m))$ and $W_{\mathbf{x}}(\mathbf{x}, m) = (W_{x^{i,j}}(\mathbf{x}, m))$ and $G_{\mathbf{x}}(\mathbf{x}, m) = (G_{x^{i,j}}^k(\mathbf{x}, m))$ 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 equation (47) is given by

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

In practice, we must solve this problem numerically. It can be written as a finite-time constrained optimization problem by finding a date, T, sufficiently far in the future that all aggregates are constant from then on. We set $T = 1000$ and solve for the discounted utilities for every person who lives until T, under the assumption that old people are non-workers from T onward. We then maximize the integral in Equation 29 from zero to T with these continuation utilities as terminal values, discounted to $t = 0$. This is subject to the laws of motion for state variables and the relevant definitions of the shares of people currently working and consumption of workers and non-workers. We implement this optimization using the OpenOCL toolbox in Matlab (Koenemann, et al. A1).

While the optimal control approach did not achieve large gains relative to our simple rule starting from April 12, we do find value in flexibility if the policy maker must commit to a path of mitigation on February 15, before infections have started to increase. We illustrate this in Figure 14. A policy maker that must commit to a policy in our parametric class cannot suppress the virus without suffering months of low consumption, whereas a policy maker with flexibility chooses to ramp up mitigation as the virus starts to spread. Therefore, restricting to our baseline functional form cannot increase utility relative to zero mitigation forever, whereas the non-parametric utilitarian optimal increases welfare for basic, luxury, and old people by 0.40%, 0.19%, and 3.88%, respectively.

Table 9: Millions of People in Each Health State

	S	A	F	E	R	D × 1000
03/31/20	318.36	2.97	2.33	0.07	6.28	5.59
04/30/20	303.11	2.57	2.53	0.13	21.60	53.38
06/29/20	249.42	1.68	1.72	0.09	46.86	154.81
09/30/20	201.42	4.31	4.59	0.24	119.03	406.81
12/31/20	171.52	0.47	0.62	0.04	156.74	599.38
12/31/21	168.82	0.00	0.00	0.00	160.56	621.95

D   Appendix References

[A1] Koenemann, J., Licitra, G., Alp, M., Diehl, M. (2017), "OpenOCL–Open Optimal Control Library".
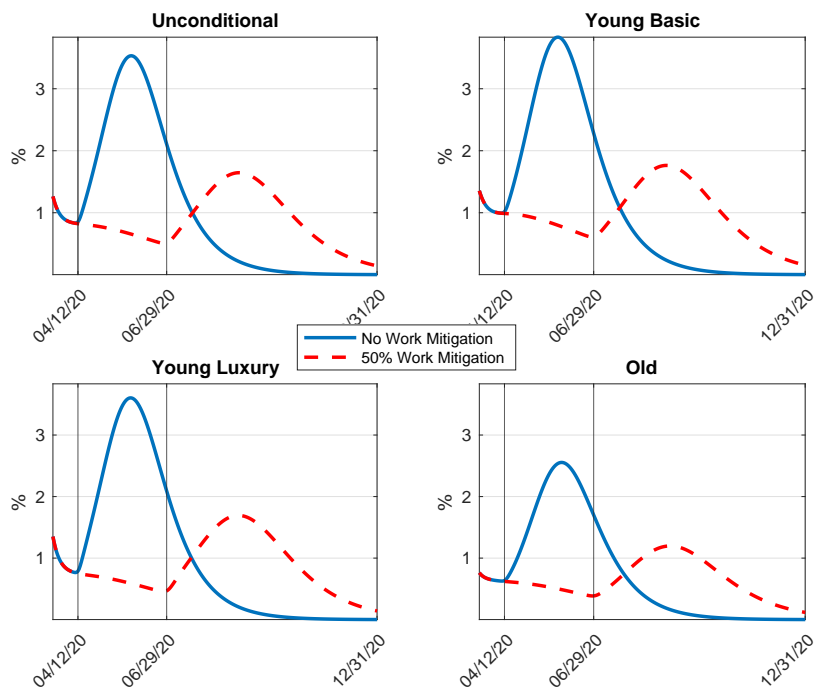
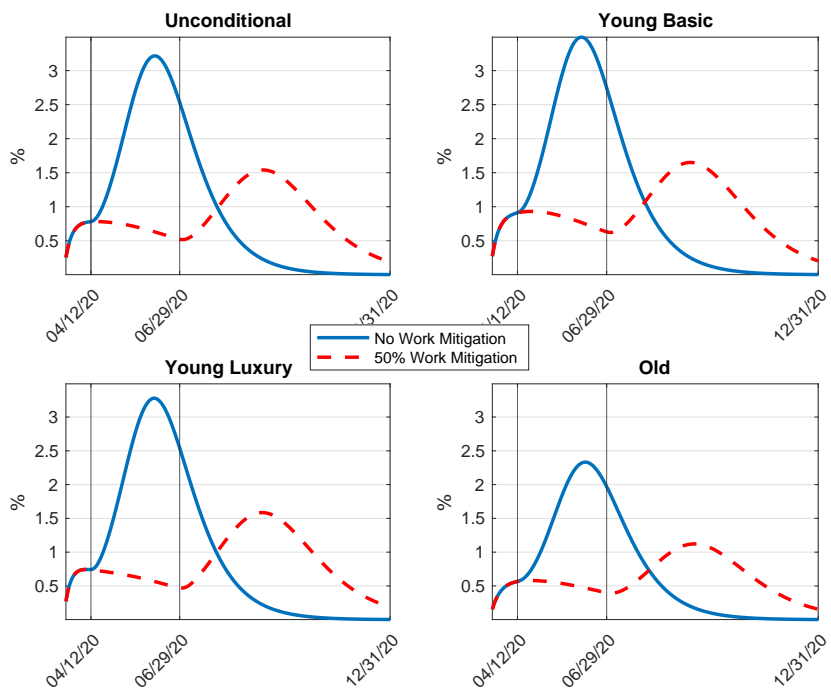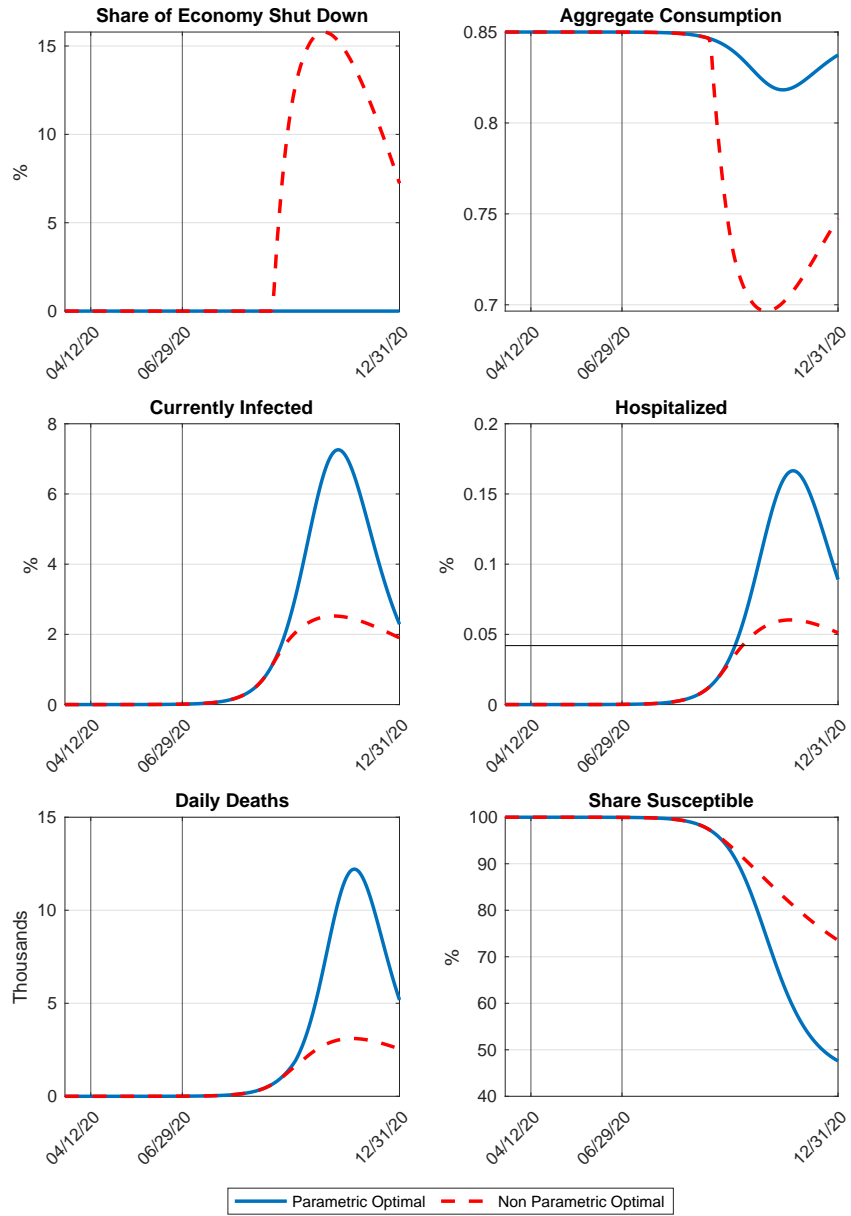Figure 12: Share of Each Group Infected but Asymptomatic

Figure 13: Share of Each Group with Fever

Figure 14: Key Outcomes, Starting Feb. 15