

Journal Pre-proof

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

Andrew Glover, Jonathan Heathcote, Dirk Krueger,
José-Víctor Ríos-Rull

PII: S0304-3932(23)00086-7
DOI: <https://doi.org/10.1016/j.jmoneco.2023.07.003>
Reference: MONEC 3507

To appear in: *Journal of Monetary Economics*

Received date: 21 December 2021
Revised date: 7 July 2023
Accepted date: 7 July 2023

Please cite this article as: Andrew Glover, Jonathan Heathcote, Dirk Krueger, José-Víctor Ríos-Rull, Health versus Wealth: On the Distributional Effects of Controlling a Pandemic, *Journal of Monetary Economics* (2023), doi: <https://doi.org/10.1016/j.jmoneco.2023.07.003>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2023 Published by Elsevier B.V.



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

Andrew Glover[†]

Jonathan Heathcote[‡]

Dirk Krueger[§]

José-Víctor Ríos-Rull[¶]

Abstract

To slow the COVID-19 virus, many countries shut down parts of the economy. Older individuals have the most to gain from slowing virus diffusion. Younger workers in shuttered sectors have most to lose. We build a model in which economic activity and disease progression are jointly determined. Individuals differ by age, sector and health status. Disease transmission occurs at work, at home, through consumption, and in hospitals. Optimal economic shutdowns in 2020 are milder when taxes are distortionary, and when the government does not have access to debt. A harder, shorter shutdown is preferred when vaccines become available in early 2021.

*We thank George Alessandria, Andrew Atkeson, Tim Lee, Simon Mongey, Andreas Peichl, Josep Pijoan-Mas, Cezar Santos, Stefanie Stantcheva, Marc Suchard as well as several anonymous referees for very helpful comments, and attendees of virtual seminars at many institutions for valuable feedback. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Banks of Kansas City and Minneapolis or the Federal Reserve System. Krueger thanks the NSF (grant SES-175708.) for financial support.

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

[‡]**Corresponding Author:** Federal Reserve Bank of Minneapolis and CEPR; Research Department, Federal Reserve Bank of Minneapolis, 90 Hennepin Avenue, Minneapolis, MN 55480 (612) 204 6385, heathcote@minneapolisfed.org

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

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

Keywords: COVID-19; Shutdowns; Redistribution; Economic Policy

J.E.L. classification codes: E1, I14, H12

Journal Pre-proof

Research Highlights

1. We explore the optimal shutdown policy in response to the COVID-19 pandemic.
2. Individuals in our quantitative model differ by age and by sector.
3. Redistribution through taxes, transfers, and debt is modeled explicitly.
4. Heterogeneity and costly redistribution significantly impact the optimal shutdown policy.
5. With vaccines a harder but short shutdown is preferred.

Journal Pre-proof

1 Introduction

A pandemic such as the COVID-19 crisis constitutes a large shock to global welfare, with adverse impacts on societal health and economic wealth. What is the optimal policy response to such shocks when social contact is central to both disease transmission and economic activity? Debate in 2020 centered on the question of how aggressively to restrict economic activity in order to slow the spread of a pandemic and how quickly to lift restrictions as it shows signs of subsiding, either naturally or in response to a vaccination campaign.

There is substantial disagreement about the answer to this question and about the factors, both in terms of the medical nature of the disease as well as the structure of the economy, that determine this answer. In this paper, we argue one source of disagreement is the fact that the benefits and costs of “lock-down” policies are large and very unequally distributed among different groups of the population. The young and the old and workers in sectors differentially impacted by lockdowns have vastly diverging preferences concerning these policies.

Standard epidemiological models miss this disagreement because they assume a representative agent setting in which all households face the same trade-off between restrictions on social interactions that slow disease transmission but also depress economic activity. In reality, for a pandemic such as COVID-19, the benefits of slower viral transmission accrue disproportionately to older households, which face a much higher risk of serious illness or death from infection. In contrast, the costs of reduced economic activity are disproportionately borne by younger households facing the brunt of lower employment. For these younger households, the costs of mitigation policies depend on their sector of work. Sensible lock-down policies designed to reduce viral spread 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. We will call this part of the economy the “luxury sector” henceforth. During the COVID-19 pandemic, for example, restaurants, bars and other establishments in the broader hospitality sector were closed first. The fact that workers cannot easily reallocate across sectors implies that lock-downs have very disparate impacts on young households specialized in different sectors. Thus, different groups in the economy (old versus young, workers in different sectors, healthy versus sick) 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 tax and

transfer policies to mitigate the increase in economic inequality that shutdowns entail. However, redistribution is costly in practice. And the more costly it is, the larger and more unequal will be the economic costs of mitigation measures. It is therefore important to study optimal lockdown and redistribution policies jointly. This is what we set out to do in this paper.

To do so, we build a novel macro-epidemiological model of a health pandemic that incorporates the interaction between macro-mitigation and micro-redistribution policies. We then apply the model to study the optimal policy response to the COVID-19 crisis, both for the first phase of the pandemic in 2020, in which we assume no vaccines were on the horizon, and then for the second phase, starting in 2021, when a gradual roll out of effective vaccines took place.¹

Our model has three key elements: (i) a household sector with heterogeneous individuals, (ii) an epidemiological block where consumption, production, and purely social interactions determine health transitions, and (iii) a government that can use distortionary taxes, transfers and debt to spread the economic costs of shutdowns across individuals and over time.

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. Consumption of basic sector output does not contribute to virus transmission, and workers in this sector are not subject to shutdowns. In contrast, the policy maker can choose to reduce employment and output in the social contact-intensive luxury sector in order to reduce virus transmission during production and social consumption.

The epidemiological model builds on a standard Susceptible-Infectious-Recovered (SIR) diffusion framework but permits a richer set of health states that are quantitatively important for our analysis. We label our variant of the SIR model the *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:

¹Our analysis of these two phases assume that policymakers do not foresee the emergence of new more contagious variants such as “Delta,” which surged in the United States in late summer 2021, or “Omicron,” which became the predominant strain in early 2022, which we think of as the third phase of the pandemic. Relatedly, they do not anticipate that in this third phase individuals who have experienced infection or been vaccinated can be subsequently reinfected. In this paper we focus on the first two phases of the pandemic. See Glover et al. (2022) for a partial analysis of the third phase with a focus on the Delta variant.

the old are much more likely to experience adverse health outcomes conditional on becoming 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 fellow 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. The three 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 infect only family members, while those in hospital care may pass the virus to health care workers.

The government maximizes a utilitarian social welfare function and has at its disposal two sets of policy levers. First, at each date it can choose what fraction of activity in the luxury sector to shut down, which we call the “extent of mitigation”. Mitigation slows the spread of the virus by reducing the rate at which susceptible workers and consumers become infected, but it also reduces to zero the market income of mitigated workers. Second, the government chooses how much income to transfer to those not working, either because they are old, because they are unwell, or because their workplaces have been closed. Transfers must be ultimately financed out of taxes on workers, but in our baseline policy specification the government can use debt to smooth tax rates over time. A utilitarian government wants to redistribute, but internalizes that higher tax rates depress labor supply and output. When mitigation increases, the government optimally trades off equity versus efficiency by both increasing tax rates – implying larger efficiency costs – and by tolerating more inequality between workers and non-workers – implying less equity. We show that the distortions induced by redistribution reduce the dynamic incentives of the government for mitigation. In particular, we prove theoretically that the marginal welfare costs of mitigation are larger when redistribution is costly than when the government has access to lump-sum taxes.

We then use this model to characterize quantitatively the optimal path of mitigation, both for the first phase of the pandemic in 2020 in the absence of a vaccine and, separately, for the second phase (the first half of 2021) when effective vaccinations that protect individuals both from contracting and from spreading the disease are gradually administered. We calibrate the model to U.S. data and show that under the actual mitigation path the model captures the dynamics of COVID-19 related deaths well. We then ask what level and time path of economic lockdowns

a utilitarian government would choose and how these contrast with the preferred policies of the three different groups of the population.

We highlight four findings. First, in the absence of a vaccine (and absent the expectation of one arriving in the near future) utilitarian optimal policy locks down about 30 percent of the nonessential sector in early 2020, with a temporary relaxation during the summer months, when infections and deaths are low. This is a compromise between vastly different policy preferences of different groups. To start with, one would expect disagreement between workers in the two different sectors, since only the luxury workers are subject to lockdown risk. However, this disagreement can effectively be addressed through a redistributive tax-transfer policy in which workers share the cost of lockdowns with the unemployed through higher taxes and transfers. Disagreement across age groups is much harder to deal with, as the old receive most of the health benefits of lockdowns and pay none of the costs in terms of higher taxes. As a result the old prefer much stricter and longer lockdowns than the young. By the same token, there is much more at stake for the old than the young, in that the welfare gains from switching from the benchmark mitigation path to the utilitarian-optimal one are about 40 times larger for the old than the young, while welfare differences across the two young groups are relatively minor.²

Second, the optimal mitigation path depends on the set of fiscal instruments to which the government has access. Suppose the government can impose type-specific lump-sum taxes, and that redistribution between workers and non-workers is therefore costless, as it implicitly is in representative agent models. Shutdowns are then less costly, and as a result optimal shutdowns are more extensive. We also consider a fiscal constitution under which the government must run a balanced budget date by date to quantify the importance of access to government debt. In that scenario, shutdowns necessitate immediate increases in tax rates, which tempers the planner's appetite for mitigation. The broader policy implication is that the optimal mitigation policy in response to a pandemic is sensitive to the details of the social insurance system and to the amount of fiscal space that a country enjoys.

Third, the optimal mitigation path is also strongly influenced by the details of how the threat to health from COVID-19 varies with age. Our utility specification incorporates an additive flow utility of being alive, whose value we calibrate to replicate a standard "value of a statistical life".

²The young account for 85 percent of the population in both the model and the data. Thus, if policy is determined by majority rule rather than by a utilitarian government, then one might expect shutdowns to reflect the preferences of the young and to therefore be too modest relative to the utilitarian optimum.

Given this specification, our planner is more concerned about total years of life lost due to COVID than about mortality *per se*. Because the risk of dying from COVID is much higher for older individuals with relatively low normal residual life expectancy, the planner in our economy chooses to mitigate less than it would if mortality rates did not vary by age, as in models based on the representative agent paradigm.

Fourth, expectations about vaccine arrival and distribution are critical determinants of the optimal profile for shutdown policies. If no vaccines are on the horizon (say in Spring 2020, the first phase of the pandemic), then the government cannot strongly affect the share of the population that will eventually fall ill. In such a scenario, optimal lockdowns are relatively modest and geared mostly toward avoiding excess demand for emergency hospital beds and associated excess mortality, which we model explicitly. In contrast, with knowledge that a vaccine is coming soon (the beginning of the second phase of the pandemic in early 2021), shutdowns can eliminate rather than merely delay infections, and much harsher lockdowns are optimal. We show that a utilitarian government chooses to lock down the entire luxury sector at the start of 2021, when it knows effective vaccines will be distributed, whereas the lockdown would be only half as extensive if there were no vaccines on the horizon. But as an increasing share of the population subsequently obtains immunity via vaccination, the government re-opens the economy much more rapidly than it would without a vaccine.

1.1 Related Literature

Our paper contributes to a by now substantial literature on the interaction between pandemics and economic activity, with a focus on the current COVID-19 crisis. Important early references include Atkeson (2020), Eichenbaum et al. (2021) and Argente et al. (2021).³

We wish to stress the following contributions of our paper relative to the literature. First, the paper is one of the first COVID-19 studies to explicitly incorporate multiple age groups. We emphasize not only the enormous age-related differences in the disease burden of COVID but also the stark policy disagreements across these different groups. We share the focus on the age dimension of heterogeneity with Acemoglu et al. (2021), Boppart et al. (2022) and Brotherhood et al. (2020).

³Other notable papers in the rapidly growing economics literature on COVID-19 include Fernandez-Villaverde and Jones (2022), Greenstone and Nigam (2020), Krueger et al. (2022), Toxvaerd (2020), Farboodi et al. (2021), Guerrieri et al. (2022), Bayer and Kuhn (2020), Berger et al. (2022), Chari et al. (2021), Hall et al. (2020) and Toda (2020).

A second distinctive feature of our framework is that the economic side of our environment is modeled in an explicit structural way. Each of the key household constituencies solves a maximization problem subject to a budget constraint, and there is no fictitious representative household that pools the economic costs of shutdowns. Therefore, the optimal policy cannot be reduced to a simple trade-off between lost output versus lives saved: the distribution of consumption and hours worked, as well as the distribution of mortality, are central policy considerations. We explore optimal policy from the perspective of a Ramsey government that uses realistically blunt policy instruments to affect household behavior. In contrast, most of the extant literature has focused on the dichotomous extremes of laissez-faire equilibrium or socially optimal allocations. The model of Kaplan et al. (2020) features richer heterogeneity than ours, but their paper does not study optimal lockdown policies or optimal redistribution.

A third novel feature of our model is that it allows us to characterize the optimal fiscal response to the COVID pandemic in closed form. In particular, we derive expressions for optimal time-varying transfers to non-workers and for optimal tax rates that depend on the expected present values of the measures of workers and non-workers in the economy. These in turn depend on the share of the population that is either mitigated or sick with COVID. We also show how the optimal fiscal response changes depending on whether or not the government has access to debt finance (see also Arellano et al. 2023). Our closed-form model of optimal redistribution could be applied in other contexts. One advantage of our model is that it accommodates transfers to individuals with zero market income, in contrast to the specification used by Benabou (2002) and Heathcote et al. (2017).

Finally, our paper is one of the few that explicitly model the interaction between the deployment of vaccines and optimal mitigation (see also Gonzalez-Eiras and Niepelt, 2020, Bognanni et al., 2020, Garriga et al., 2022). Gollier (2021) also explores the positive and normative effects of the extent and timing of vaccine deployment in a model with multiple age groups. But he considers neither the policy conflicts among age groups nor the jointly optimal redistribution and lockdown policies.

In Section 2, we describe how we model the joint evolution of the economy and the population. In Section 3, we then explain how we model mitigation and redistribution policies and the optimal policy problem. The calibration strategy is discussed in Section 4. The findings are in Section 5, and Section 6 discusses optimal policy in the presence of a vaccination campaign. Section 7

concludes.

2 The Model

We model an open economy in continuous time in which the government can borrow and lend freely at a fixed exogenous interest rate. We first describe the individual state space, spelling out the nature of heterogeneity by age and health status. In Section 2.2, we describe the multi-sector production technology and explain 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. In Section 2.4 and Section 2.5 we set out the lifetime and period utility function and labor supply decisions, respectively, and Section 2.6 discusses the aggregation of the household and government sector. Section 2.7 discusses the key model assumptions we have made.

2.1 Household Heterogeneity

Time starts at $t = 0$ 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. Parameters are denoted with Greek letters. Generically, we use the letter x to denote population measures, with superscripts specifying subsets of the population.

Agents can be young or old, denoted by y and o . We think of the young as individuals below the age of 65 and their measure is given by x^y . For simplicity, and given the short time horizon of interest, we abstract from population growth and from aging and death unrelated to COVID-19 during the period of analysis.⁴ Within each age group, agents are differentiated by health status, i , 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 I group in the standard *SIR* model – and can pass it onto others. However, they differ in their symptoms. The asymptomatic have no symptoms or only 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 or only very mild symptoms. Those with a fever are sufficiently sick to know they

⁴Thus, there are no individuals who enter the economy during the pandemic; for an analysis of the differential welfare effects of aggregate shocks between newborn and older individuals, see Glover et al. (2020b).

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 (s) to asymptomatic (a) to fever (f) to emergency care (e) to dead (d).⁵ However, recovery (r) is possible from any of the a , f and e states.

2.2 Economic Activity: Technology and Mitigation

Young agents in the model are further differentiated by the sector in which they can work. A mass x^b of the young work in an essential, basic b sector, while the rest of the young of mass x^ℓ work in a luxury, non-essential sector, denoted ℓ . Shutdowns only apply to the ℓ sector, and require some or all of the ℓ sector workers to stay at home in order to reduce virus transmission in the workplace and through luxury consumption activities. We call such a policy a mitigation policy, and use m_t to denote the fraction of luxury workers who are instructed to not work at time t . We assume that individuals cannot change the sector they work in. In terms of notation, superscripts denote the dimensions of household heterogeneity, that is, 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 in the susceptible health state.

The production technology is linear in labor input in both sectors. Thus, output in the basic sector is given by the measure $x^{bw} = x^{ybs} + x^{yba} + x^{ybr}$ of young workers employed there, times the number of hours h^b they work:

$$Y^b = [x^{ybs} + x^{yba} + x^{ybr}] h^b = x^{bw} h^b. \quad (1)$$

Note that this specification assumes that asymptomatic individuals carrying the virus continue to work while those with fever stay at home.⁶ In contrast to the basic sector, output in the luxury sector depends on mitigation policy and is given by

$$Y^\ell(m_t) = (1 - m_t) [x^{y\ell s} + x^{y\ell a} + x^{y\ell r}] h^\ell = (1 - m_t) x^{\ell w} h^\ell. \quad (2)$$

⁵Note that in the standard SEIR model, agents in the exposed state E have been subjected 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. (2022) make a similar modeling choice.

⁶One could instead imagine a policy of tracing contacts of infected people, which would allow the government to keep some portion of asymptomatic workers at home.

Since production is linear in labor in both sectors with identical productivity, the relative price of luxury goods in terms of the basic good is 1, and GDP is given by $Y = Y^b + Y^\ell$.

2.3 Health Transitions: The SAFER Model

We now describe the dynamics of individuals across health states. At date $t = 0$, the total mass of living individuals is one, μ^y denotes the share that is young, and μ^b is the share of the young who work in the basic sector. At each point in time, we denote populations by age and sector by $x^{yb} = \sum_{i \in \{s, a, f, e, r\}} x^{ybi}$, $x^{y\ell} = \sum_{i \in \{s, a, f, e, r\}} x^{y\ell i}$, and $x^o = \sum_{i \in \{s, a, f, e, r\}} x^{oi}$. Thus, at $t = 0$, $x^{yb} = \mu^y \mu^b$, $x^{y\ell} = \mu^y (1 - \mu^b)$, and $x^o = (1 - \mu^y)$. At any point in time 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, 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. The number of such workers who catch the virus is their mass (x^{ybs} for young basic sector workers, for example) times the number of virus-transmitting interactions they have. 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, 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, denoted by $x_j(m_t)$ for $j \in \{w, c, h, e\}$, and the likelihood that such meetings result in infection, which we label infection-generating rates $\beta_j(m_t)$.

The numbers of contagious people in each activity are given by the following population measures

$$x_w(m_t) = x^{yba} + (1 - m_t)x^{y\ell a}, \quad (3)$$

$$x_c = x^a, \quad (4)$$

$$x_h = x^a + x^f, \quad (5)$$

$$x_e = x^e, \quad (6)$$

where these measures reflect the assumptions that symptomatic (s or e) individuals neither work nor shop and that basic and luxury sector workers can meet in the workplace. Note that the

number of contagious workers depends on the mitigation choice m_t .

In modeling the infection-generating rates, we recognize that 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. Assume workers are assigned to a unit interval of sub-sectors $i \in [0, 1]$ where sub-sectors are ranked from those generating the least social interaction to those generating the most. Also assume the sub-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 governing the intensity by which meetings generate infections. When the government asks a fraction m_t of luxury workers to stay at home, we assume it targets the sub-sectors generating the most interactions; that is, $i \in [1 - m_t, 1]$. The average infection-generating rates of the sub-sectors that remain are then $\alpha_w(1 - m_t)$ and $\alpha_c(1 - m_t)$, respectively.⁷ By assumption, the government does not mitigate any workers in the basic sector, so the average workplace infection-generating rate in that sector is α_w . The economy-wide infection-generating rate for work-related infections is the following employment-share-weighted average across the two sectors:

$$\beta_w(m_t) = \frac{x^{bw}}{x^{bw} + (1 - m_t)x^{\ell w}} \times \alpha_w + \frac{(1 - m_t)x^{\ell w}}{x^{bw} + (1 - m_t)x^{\ell w}} \times \alpha_w(1 - m_t). \quad (7)$$

The infection-generating rate for consumption, $\beta_c(m_t)$, is similar, except that (i) only luxury consumption is associated with infection risk, and (ii) we assume that the infection-generating rate is proportional to the number of luxury sector workers working, which we think of as a proxy for the number of stores that are open. Thus,

$$\beta_c(m_t) = \frac{(1 - m_t)x^{\ell w}}{(1 - \mu^b)\mu^y} \times \alpha_c(1 - m_t), \quad (8)$$

where the denominator is the pre-mitigation number of luxury workers. The key property of these expressions is that as mitigation increases, the average social-interaction-generating rate falls. When all workers are healthy and there is no mitigation, $\beta_w(0) = \alpha_w$ and $\beta_c(0) = \alpha_c$.

⁷ $E[2\alpha_w i | i \leq (1 - m_t)] = \frac{2\alpha_w}{1 - m_t} \int_0^{1 - m_t} i di = \frac{2\alpha_w}{1 - m_t} \frac{(1 - m_t)^2}{2} = \alpha_w(1 - m_t)$.

Equations (9)-(11) below capture the flows of basic sector workers, luxury sector workers, and older individuals out of the susceptible state and into the asymptomatic state:

$$\dot{x}^{ybs} = -[\beta_c(m_t)x_c + \beta_h x_h] x^{ybs} - \beta_w(m_t)x_w(m_t) x^{ybs} - \beta_e x_e x^{ybs}, \quad (9)$$

$$\dot{x}^{y\ell s} = -[\beta_c(m_t)x_c + \beta_h x_h] x^{y\ell s} - \beta_w(m_t)x_w(m_t)(1 - m_t)x^{y\ell s}, \quad (10)$$

$$\dot{x}^{os} = -[\beta_c(m_t)x_c + \beta_h x_h] x^{os}. \quad (11)$$

Consider the first outflow rate in equation (9). The flow of young basic sector workers getting infected through consumption is the number of such workers who are susceptible, x^{ybs} , times the number of contagious shoppers, x_c , times the infection-generating rate, $\beta_c(m_t)$. The flow of young basic sector workers getting infected from co-workers is similarly constructed.

The rate at which young basic workers contract the virus at home, $\beta_h x_h$, depends on the number of contagious workers in the household, x_h , defined in equation (5). Note that both asymptomatic and fever-suffering individuals are at home. We assume that caring for those requiring emergency care is a task that falls entirely on basic workers. The risk of contracting the virus from this activity is proportional to the number of hospitalized people, $x_e = x^e$, with infection-generating rate β_e , which reflects the strength of precautions taken in hospitals.

Parallel to equation (9), equation (10) 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 work reduced hours when $m_t > 0$ and thus have fewer work interactions in which they could get infected. Furthermore, luxury sector workers do not take care of hospital patients, and thus the last term in equation (9) is absent in equation (10). Equation (11) displays infections among the old who get infected only from market consumption and from interactions at home.

The remainder of the epidemiological block simply describes the transition of individuals through the health states (asymptomatic, fever-suffering, hospitalized, and recovered) once they have been infected. These transitions are described in equations (26) to (37) 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 $\sigma^{yed} + \varphi(x^e)$ and $\sigma^{oed} + \varphi(x^e)$, where φ is the excess mortality

rate when hospital capacity is overused and is given by

$$\varphi(x^e) = \lambda_o \max\{x^e - \Theta, 0\}. \quad (12)$$

In (12) the term in the max operator defines the extent of hospital overuse given capacity Θ , treated as fixed in the time horizon under consideration. 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 Lifetime Utility Function

Preferences are defined over consumption and hours worked and also incorporate utility from being alive and being in a specific health state. Lifetime utility for the young is given by

$$E_0 \left\{ \int_0^{T^y} e^{-\rho t} S_t^y [u(c_t^y, h_t^y) + \bar{u} + \hat{u}_t^i] dt \right\}, \quad (13)$$

where ρ denotes the discount rate, T^y is remaining life expectancy (absent premature death from COVID), and S_t^y denotes the probability of surviving to date t . Flow utility at date t is the sum of a term involving consumption and labor supply, $u(c_t^y, h_t^y)$, a flow utility from simply being alive, \bar{u} , and a flow value that varies with health state i , \hat{u}_t^i . We assume that $\hat{u}_t^s = \hat{u}_t^a = \hat{u}_t^r = 0$ and that $\hat{u}_t^e < \hat{u}_t^f < 0$. Thus, having a fever is bad, and being treated in the hospital is worse. If an individual dies of COVID, all utility terms are zero thereafter. Preferences for the old are similar, except that they do not work, so $h_t^o = 0$. In addition, the old have a shorter normal residual life expectancy, T^o , and face greater COVID mortality risk, reflected in lower survival probabilities, S_t^o .

In equilibrium, expected utility of a young individual will depend on the sector in which she works, for two reasons. 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.

2.5 The Period Utility Function, Household Consumption and Labor Supply

Households value consumption and labor supply (if they work) according to the following Greenwood, Hercowitz and Huffman style utility function:

$$u(c, h) = \log\left(c - \frac{h^{1+\frac{1}{\chi}}}{1 + \frac{1}{\chi}}\right),$$

where utility from household consumption $c = c^b + c^\ell$ is a linear aggregate of consumption in the two sectors. Since one unit of labor produces one unit of output in both sectors, the wage of in both sectors is $w^b = w^\ell = 1$. Recall that the sector an individual can work in is part of her type and cannot be adjusted during the time horizon under consideration.

The government taxes labor income at a flat rate τ and provides everybody not working with a transfer T , which is simply consumed. Healthy, non-mitigated households solve

$$\begin{aligned} \max_{c, h} U &= \log\left(c - \frac{h^{1+\frac{1}{\chi}}}{1 + \frac{1}{\chi}}\right), \\ \text{s.t. } c &= (1 - \tau)h, \end{aligned}$$

with solution

$$h = (1 - \tau)^\chi, \quad (14)$$

$$c = (1 - \tau)h = (1 - \tau)^{1+\chi}, \quad (15)$$

$$U = -\log(1 + \chi) + (1 + \chi)\log(1 - \tau). \quad (16)$$

For non-working households, the budget constraint and period utility are

$$c^n = T, \quad (17)$$

$$U^n = \log(c^n) = \log(T). \quad (18)$$

2.6 The Government and Aggregation

The government purchases goods of both sectors of the economy and raises taxes at rate τ to pay for these expenditures and for transfers T . We assume that government purchases are a

constant share g of domestically produced output in both sectors: $G^i = gY^i$ for $i \in \{b, \ell\}$. For future reference, we write the measures of working, nonworking and total individuals as

$$\begin{aligned} x^w(m) &= x^{bs} + x^{ba} + x^{br} + (1 - m)(x^{\ell s} + x^{\ell a} + x^{\ell r}) = x^{bw} + (1 - m)x^{\ell w}, \\ x^n(m) &= x^o + x^{bf} + x^{be} + x^{\ell f} + x^{\ell e} + m(x^{\ell s} + x^{\ell a} + x^{\ell r}). \\ x(m) &= x^w(m) + x^n(m). \end{aligned}$$

2.7 Discussion of Model Assumptions

Our economy has the feature that the economic side of the model is analytically tractable. As we will shortly see, this makes it possible to jointly characterize optimal mitigation and redistribution policies. However, our simple model abstracts from several ingredients that may be important for modeling the COVID-19 pandemic.

First, and perhaps most importantly, individuals in our model cannot take any choices that directly reduce their infection risk.⁸ Eichenbaum et al. (2021), Farboodi et al. (2021) and Engle et al. (2021) are important examples of models that endogenize behavioral responses to mitigate infection risk. If agents in our model could reduce their risk of infection by reducing hours of work or consumption, economic activity would endogenously contract when the pandemic hits, and there would be less need for government intervention to further restrict economic activity.

However, even if we were to give individuals scope to reduce their infection risk, government restrictions on activity would still be optimal, for two reasons. The first reason is the standard externality logic: when trading off the costs of reduced economic activity against the benefits of reduced infection risk, selfish individuals internalize the payoffs from reducing own infection risk, but not the benefits of reducing the risk of infecting others. The second reason we would not expect adequate endogenous private mitigation is that in our heterogeneous agent model, most virus transmission occurs among young individuals. For the young, the risk of dying from COVID is very small, and thus young individuals in the model will not choose to take very costly actions to reduce their infection risk. Thus, if we were to introduce it, endogenous private mitigation would be much more limited in our heterogeneous agent model than in a representative agent

⁸Our calibration captures these choices in a simple reduced-form way by imposing a one-time discrete decline in infection-generating rates when COVID first emerges, and a gradual increase in those rates in 2021 that we link to the pace of vaccinations.

setting.

In addition, with or without endogenous private mitigation, optimal policy trades off the marginal costs of additional economic restrictions against the marginal benefits of better health outcomes. Thus, given optimal policy, we would expect an extended model with private mitigation to exhibit very similar or identical dynamics for health and economic outcomes to the ones we compute. Only the decomposition between the portion of the economic slowdown that is endogenous versus policy-directed would change.

We have considered one natural extension of the model to introduce a behavioral margin by allowing individuals to choose to not work –i.e., to volunteer to be mitigated– and to thereby eliminate the risk of infection in the workplace. In this extension, workers who choose to not work receive transfers. We find that, given our utilitarian optimal path for mitigation, neither basic nor luxury sector workers would ever choose not to work.⁹ The logic is simply that the private cost of being mitigated – in terms of lost current income – outweigh the private gains, in terms of reduced infection risk and extended longevity.¹⁰

Note that in this extended model, public and private mitigation measures are equally efficient, so there is no motive for the government to moderate government shutdowns in order to encourage more workers to voluntarily stay at home.¹¹ Thus, our utilitarian optimal mitigation path would

⁹The specific thought experiment is as follows. Given perfect foresight over the aggregate path of the pandemic, we evaluate expected lifetime utility at the start of the pandemic for young individuals in our baseline economy and for hypothetical young individuals who know they will not work for the duration of one future calendar month – for example, December 2020.

¹⁰To understand this, consider the following calculation that gives a sense of willingness-to-pay. Through December 1 2021, the fraction of Americans below the age of 65 who experienced COVID-related deaths was 0.072 percent (the CDC reports 195,195 such deaths). Given our calibrated values for ρ , T^y and \bar{u} we can ask how much consumption a young individual would be willing to forgo, for one year, to completely eliminate the risk of dying from COVID. Given consumption and hours equal to the per capita average values \bar{c} and \bar{h} , the answer is the value for x that solves

$$\begin{aligned} & \int_0^{365} \exp(-\rho t) \left(\log \left((1-x)\bar{c} - \frac{\bar{h}^{1+\frac{1}{x}}}{1+\frac{1}{x}} \right) + \bar{u} \right) dt + \int_{365}^{T^y} \exp(-\rho t) \left(\log \left(\bar{c} - \frac{\bar{h}^{1+\frac{1}{x}}}{1+\frac{1}{x}} \right) + \bar{u} \right) dt \\ &= (1 - 0.01 * 0.072) \int_0^{T^y} \exp(-\rho t) \left(\log \left(\bar{c} - \frac{\bar{h}^{1+\frac{1}{x}}}{1+\frac{1}{x}} \right) + \bar{u} \right) dt \end{aligned}$$

which is 10.6 percent of average annual consumption, or \$4,675. This number is small relative to the income loss from not working, even when the utility benefit from the reduction in hours worked is taken into account.

¹¹In this version of the model the government uses distortionary taxes to fund transfers to all non-workers, irrespective of whether transfer recipients have been told to not work by the government or have chosen to not work to reduce infection risk.

be unchanged if people could choose not to work.¹²

Two other model assumptions that are important for retaining tractability on the economic side of the model are that households cannot save, and that basic and luxury goods enter separably in preferences. See Kaplan et al. (2020) and Guerrieri et al. (2022) for examples of papers that allow for savings and that introduce richer sectoral variation and cross-sector spillovers.

On the health side, we refine the textbook SIR model by splitting the infected state into three sub-states, the asymptomatic A , fever F , and emergency room E states. Our motivation for doing so is two-fold. First, we think a quantitative analysis of optimal lockdowns needs to distinguish between infected individuals without and with symptoms. We find it plausible to assume that individuals with COVID symptoms (those in the F state) will stay home, and not work or shop. However, individuals infected with COVID are typically contagious for several days before they develop symptoms, and it is the presence of these individuals in the A state that creates a rationale for lockdowns. In particular, the value of lockdowns depends crucially on the relative shares of A versus F individuals in the population: lockdowns would be pointless in our model if every infected person was symptomatic (F state), and would be most valuable if the infected were all asymptomatic (A state). Thus, getting these shares approximately correct (which the *SAFER* model allows) seems crucial for quantifying optimal policy.

Second, the presence of the E state allows us to model the potential for hospital overload (especially in the first phase of the pandemic). Section 5.1 demonstrates that including this overload mechanism in the model is quantitatively very important for understanding why the optimal mitigation path seeks to “flatten the curve.”

3 Fiscal Policy: Taxes, Transfers and Government Debt

As described above, in addition to financing exogenous public spending parameterized by g , the government is responsible for three choices: a path of mitigation (shutdowns) m_t , redistribution through proportional taxation on workers at rate τ_t and lump-sum transfers T_t to individuals who do not or cannot work (which include those unemployed because of shutdowns, those with fever or who are hospitalized, and those who have retired). These fiscal choices imply a path of

¹²One could imagine alternative models in which private mitigation is more efficient than public mitigation. For example, suppose COVID mortality risk varies across workers, and that individual workers are better informed about their idiosyncratic risk than the government. In such an environment a better policy than mitigating workers at random might be to let individuals decide whether or not to work, while providing transfers to those who choose not to.

government debt $B = (B_t)_{t=0}^{\infty}$, given an initial level B_0 . We assume that the government can borrow at a fixed exogenous interest rate, assumed to equal the household discount rate ρ , and that when setting fiscal policy, it values all individuals, implying equal Pareto weights in the social welfare function.

Government spending, taxes and transfers are restricted by the following intertemporal budget constraint (with the path of government debt ensuring that the instantaneous budget constraint is satisfied). To interpret this constraint, recall that \bar{T} marks the end of the COVID-19 planning horizon, after which there are no further COVID-19 infections and thus no need for mitigation efforts. The planning horizon of the government ends at time T^o , when all transfer-receiving old individuals die.¹³

$$\int_0^{\bar{T}} e^{-\rho t} ((g - \tau_t)(1 - \tau_t)^{\chi} x_t^w(m_t) + x_t^n(m_t)T_t) dt + \int_{\bar{T}}^{T^o} e^{-\rho t} ((g - \tau_t)(1 - \tau_t)^{\chi} x_t^w(0) + x_t^n(0)T_t) dt = 0 \quad (19)$$

Government debt during the COVID period evolves according to

$$\dot{B}_t = (g - \tau_t)(1 - \tau_t)^{\chi} x_t^w(m_t) + x_t^n(m_t)T_t + \rho B_t$$

with initial condition $B_0 = 0$ and assumed terminal condition $B_{T^o} = 0$, so that all debt accumulated during the COVID-19 period is repaid during the remaining lifetime of the old.

3.1 Optimal Fiscal Policy for a Given Mitigation Path

Since fiscal policy does not affect health transitions, we can solve the optimal Ramsey policy problem in two steps. In a first step, given a path of mitigation $\mathbf{m} = (m_t)$ and the associated path for the health distribution \mathbf{x} , the government chooses fiscal policy to maximize the economic part of social welfare $\mathcal{W}(\mathbf{x}, \mathbf{m})$, ignoring the additive health-related terms:

$$\mathcal{W}(\mathbf{x}, \mathbf{m}) = \max_{\tau_t, T_t} \left\{ \int_0^{\bar{T}} e^{-\rho t} \left[x_t^w(m_t) \log \left(\frac{(1 - \tau_t)^{1+\chi}}{1 + \chi} \right) + x_t^n(m_t) \log(T_t) \right] dt \right. \quad (20)$$

$$\left. + \int_{\bar{T}}^{T^o} e^{-\rho t} \left[x_t^w(0) \log \left(\frac{(1 - \tau_t)^{1+\chi}}{1 + \chi} \right) + x_t^n(0) \log(T_t) \right] dt \right\} \quad (21)$$

subject to the intertemporal government budget constraint, eq. (19).

¹³Even though the young live for another $T^y - T^o$ periods (something the social welfare function will capture), the fiscal side of the model ends at time T^o with the departure of the old.

For a given path of mitigation, the following proposition completely characterizes the optimal fiscal policy chosen by the government. The proof is a standard characterization through first order conditions with respect to taxes and transfers (τ_t, T_t) at each instant, and is provided in Appendix B.1

Proposition 1 *For a given path (m_t) , the optimal paths of tax rates and transfers (τ_t, T_t) are constant and given by*

$$\tau^*(\Theta) = g + \frac{(1-g)}{(1+\chi)(1+\Theta)}, \quad (22)$$

$$T^*(\Theta) = \Theta(\tau^*(\Theta) - g)(1 - \tau^*(\Theta))^\chi = \frac{\Theta}{1+\Theta} \left(\frac{1-g}{1+\chi} \right)^{1+\chi} \left(1 + \chi - \frac{1}{1+\Theta} \right)^\chi, \quad (23)$$

$$\text{where } \Theta = \frac{\int_0^{T^o} e^{-\rho t} x_t^W(m_t) dt}{\int_0^{T^o} e^{-\rho t} x_t^N(m_t) dt} = \frac{W}{N} \quad (24)$$

is a summary measure of the impact of mitigation on the ratio of workers W to non-workers N in the economy (where W and N are appropriately discounted present values). The optimal tax rate $\tau^*(\Theta)$ is strictly decreasing in Θ (and thus strictly increasing in mitigation) and strictly increasing in spending g .

In order to highlight the role of costly redistribution we will contrast our baseline model with an alternative fiscal constitution under which the government can raise funds for redistribution costlessly. In this alternative economy, the government will optimally equalize period utilities from consumption-labor allocations between workers and non-workers, and thus the economy with costless redistribution behaves exactly like a representative agent model from an economic standpoint.

More specifically, we now assume the government has access to lump-sum taxes T_t^W on workers and transfers T_t to non-workers, but is otherwise subject to the same intertemporal government budget constraint and maximizes the same objective as in the case of distortionary taxation.¹⁴ The following proposition characterizes the economic equilibrium and optimal fiscal policy, given a path for mitigation.

¹⁴In general the lump-sum tax could differ across workers in different sectors, and the transfer could differ between mitigated versus sick workers and retired individuals. But given a utilitarian objective, this sort of variation is not optimal.

Proposition 2 *In the lump-sum tax economy (where we index taxes by superscript (w,L) and transfers by superscript L) labor supply, consumption, and period economic welfare for workers are*

$$h = 1, c = 1 - T_t^{w,L}, U = \log\left(\frac{1}{1+\chi} - T_t^{w,L}\right)$$

while the corresponding allocations for non-workers are

$$h^n = 0, c^n = T_t^L, U^n = \log(T_t^L).$$

Optimal lump-sum taxes and transfers are constant and given by

$$\begin{aligned} T^{*(w,L)}(\Theta) &= g + \frac{(1-g)}{(1+\chi)(1+\Theta)} - \frac{g\chi}{(1+\chi)(1+\Theta)} \\ T^{*L}(\Theta) &= \frac{\Theta}{1+\Theta} \left(\frac{1-g(1+\chi)}{1+\chi} \right) \end{aligned}$$

and there is no utility inequality in the economy:

$$U^i = \log\left(\frac{1}{1+\chi} - T^{*(w,L)}\right) = \log(T^{*L}) = U^n$$

Since labor supply in the lump-sum economy is equal to one, we can interpret $T^{w,L}(\Theta)$ as a tax rate.¹⁵

In our quantitative analysis we will assess the importance of the government's ability to smooth the economic costs of mitigation through the use of government debt. To that end, we will compare the outcomes in the benchmark economy with government debt to a world in which the government has no fiscal space and needs to run a balanced budget in every period. The previous propositions about optimal fiscal policy given mitigation straightforwardly generalize to such a balanced budget world, and are summarized in the following corollary whose proof is contained in Appendix B.4.

¹⁵In the absence of labor supply distortions (i.e., with $\chi = 0$) we have

$$\begin{aligned} \tau^*(\Theta) &= g + \frac{1-g}{1+\Theta} = T^{*(w,L)}(\Theta) \\ T^*(\Theta) &= \frac{\Theta}{1+\Theta} (1-g) = T^{*L}(\Theta) \end{aligned}$$

and allocations and optimal policy coincide in the two economies, given the same path for mitigation.

Corollary 1 *Given a path for mitigation (m_t), a government facing a balanced budget requirement sets time-varying taxes and transfers according to*

$$\begin{aligned}\tau_t^* &= g + \frac{(1-g)}{(1+\chi)(1+\Theta_t)} \\ T_t^* &= \frac{\Theta_t}{1+\Theta_t} \left(\frac{1-g}{1+\chi} \right)^{1+\chi} \left(1+\chi - \frac{1}{1+\Theta_t} \right)^\chi.\end{aligned}$$

in the economy with distortionary taxation, and

$$\begin{aligned}T_t^{w,L} &= g + \frac{(1-g)}{(1+\chi)(1+\Theta_t)} - \frac{g\chi}{(1+\chi)(1+\Theta_t)} \\ T_t^L &= \frac{\Theta_t}{1+\Theta_t} \left(\frac{1-g(1+\chi)}{1+\chi} \right)\end{aligned}$$

*in the economy with lump-sum taxation. Here $\Theta_t = \frac{x_t^w(m_t)}{x_t^n(m_t)}$ is the ratio of **current** workers to non-workers, and is strictly decreasing in mitigation m_t . That is, optimal tax rates increase in response to an increase in current mitigation m_t .*

3.2 The Impact of Mitigation on Economic Welfare

The full characterization of allocations and optimal fiscal policy allows us to evaluate how the economic part of social welfare depends on mitigation m_t . The direct economic impact of a marginal increase in mitigation m_t at instant t is that it reduces the summary measure of workers W by $-e^{-\rho t} x_t^{\ell w}$ and increases the summary measure N of non-workers by the same amount. Of course, a change in m_t changes, through the health part of the model, (x_t^w, x_t^n) for future time periods τ , and thus it changes social welfare through the health part of the model. We now characterize theoretically the direct (or static) economic impact of m_t , ignoring the indirect impact on social welfare \mathcal{W} through health transitions and the associated change in the population health distribution. Denoting this direct impact by $\frac{\partial \mathcal{W}^D}{\partial m_t}$ and $\frac{\partial \mathcal{W}^L}{\partial m_t}$ for the distortionary and the lump-sum economy, respectively, we have the following proposition, proved in Appendix B.3.

Proposition 3 *The direct economic welfare impact of mitigation is negative, more strongly so in the economy with tax distortions:*

$$\frac{\partial \mathcal{W}^L}{\partial m_t} = -e^{-\rho t} x_t^{\ell w} \frac{\partial \mathcal{W}^L}{\partial W} = -e^{-\rho t} x_t^{\ell w} \frac{W+N}{W} < 0$$

$$\frac{\partial \mathcal{W}^D}{\partial m_t} = \frac{\partial \mathcal{W}^L}{\partial m_t} - e^{-\rho t} x_t^{\ell w} \log \left(1 + \left(\frac{\chi}{1+\chi} \right) \left(\frac{W+N}{W} - 1 \right) \right) < \frac{\partial \mathcal{W}^L}{\partial m_t}$$

The same ranking of marginal economic welfare costs of mitigation obtains if we compare the two tax systems in the absence of government debt.

The key intuition from the propositions is that more mitigation means higher optimal tax rates. These higher tax rates translate into larger distortions, which amplify the negative effect of mitigation m on economic welfare, relative to a hypothetical alternative economy in which resources can be redistributed to non-working households in a non-distortionary fashion. The presence of government debt allows the government to smooth the cost of mitigation (and in the case of distortionary taxation, the time-varying distortions) over time. These costs have to be traded off against the health benefits of mitigation in the form of a more favorable evolution of the population health distribution. The optimal mitigation policy characterized quantitatively in the remainder of the paper resolves this trade-off.

3.3 Optimal Policy

The dynamic problem of the government is to choose an optimal path of mitigation $m(t)$ and associated health distribution \mathbf{x}_t to maximize

$$\max_{(\mathbf{x}, \mathbf{m})} \left\{ \mathcal{W}(\mathbf{x}, \mathbf{m}) + \int_0^{\bar{T}} e^{-\rho t} \left[x_t \bar{u} + \sum_{j \in \{yb, y\ell, o\}, i \in \{s, a, f, e, r\}} x_t^{ji} \widehat{u}^i \right] dt + \int_{\bar{T}}^{T^o} e^{-\rho t} x_t^o \bar{u} dt + \int_{\bar{T}}^{T^y} e^{-\rho t} x_t^y \bar{u} dt \right\} \quad (25)$$

The maximization is subject to the law of motion of the health population distribution $\mathbf{x} = (x_t)$. The objective of the government has three components: economic welfare under optimal fiscal policy welfare $\mathcal{W}(\mathbf{x}, \mathbf{m})$ (reflecting welfare from consumption-labor allocations), the value of life \bar{u} for everyone alive during the COVID-19 period $[0, \bar{T}]$ and during the remaining lifetime of the old and young survivors of COVID-19, and finally the health discounts \widehat{u}^i of those in health states i .

The optimal policy path is the solution to this optimal control problem. Its key trade-off is that a marginal increase in mitigation m entails static economic costs of $\frac{\partial \mathcal{W}(\mathbf{x}, \mathbf{m})}{\partial m_t}$, characterized in Proposition 3, stemming from an increase in the tax rate. The dynamic benefit is a more favorable change in the population health distribution: an increase in m reduces the outflow of

individuals from the susceptible to the asymptomatic state.

4 Calibration

4.1 Time-Invariant Parameters

We set the population share of the young, μ^y , to 85 percent, which is the current fraction of the U.S. population below the age of 65.

Preferences The pure time discount rate is assumed to be 3 percent on an annual level. We set the Frisch labor supply elasticity χ to one.

We set the value of a statistical life year (VSLY) to 6.25 times yearly per capita consumption in 2019, which implies a VSLY of \$276,700.¹⁶ This is very similar to the value used by Hall et al. (2020). The average age of Americans is 37.9, and at that age, remaining life expectancy is 42.6 years.¹⁷ If we discount at 3 percent, a flow payment of \$276,700 for 42.6 years is equivalent to a one-time payment of \$6.65 million (the value of a statistical life).

To translate this VSLY into a value for \bar{u} we conduct the following standard thought experiment. Suppose there is some policy intervention which (across a large population) saves one year of life for a “representative” individual with average consumption \bar{c} and working average hours \bar{h} . Households will be indifferent about this intervention when the value of the extra year of life is equal to the value of a statistical life year, i.e., when

$$\underbrace{\log\left(\bar{c} - \frac{\bar{h}^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}}\right) + \bar{u}}_{\text{utility value of a year of life}} = \underbrace{\frac{1}{\bar{c} - \frac{\bar{h}^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}}}}_{\text{marginal utility}} \times \underbrace{6.25\bar{c}}_{\text{VSLY}},$$

where the VSLY is multiplied by the marginal utility of consumption to convert the VSLY into utility units. The implied value for \bar{u} is 11.61.¹⁸

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

¹⁶Per capita consumption in 2019 was \$44,272 : see <https://fred.stlouisfed.org/series/A794RCOA052NBEA>.

¹⁷We take life expectancy values from Table A of Arias and Xu (2020).

¹⁸Note that many calculations for the value of a statistical life abstract from labor supply.

32.5 and 72.5 years old, with corresponding expected residual life expectancy of $T^y = 47.8$ and $T^o = 14.0$ years. These values imply remaining values of life at the start of the pandemic of $(1 - e^{-\rho T^y})/\rho = 25.4$ times \bar{u} and $(1 - e^{-\rho T^o})/\rho = 11.4$ times \bar{u} , respectively. Thus, for a utilitarian government, each COVID death of a young individual in the model will be roughly twice as costly as a death of an old individual.

Sectors In our model the differences between the basic and luxury sectors of the economy are that (1) only luxury consumption is a source of COVID infection, and (2) shutdowns are assumed to be concentrated in the luxury sector. To calibrate the relative employment and output shares of the two sectors, we partition the components of consumer expenditure in the CPI into those that have a social aspect of consumption and those that do not.¹⁹ This partition implies that pre-COVID, the basic sector accounts for $\mu^b = 55$ percent of the economy.

We set government's share of output (symmetric across sectors) to $g = 0.247$, which corresponds to government outlays as a share of GDP in 2019. We exclude Social Security and Medicare from this measure of spending, since we think of those as part of model transfers.²⁰

Given $g = 0.247$ and $\chi = 1.0$, the optimal pre-COVID tax rate given by equation 22 is $\tau^* = 0.303$, implying an output level of $Y = \bar{h} = \mu^y(1 - \tau^*) = 0.592$, and consumption levels for workers and for the old of $c^y = (1 - \tau^*)^2 = 0.485$ and $c^o = T^* = 0.223$.

Disease Progression There are twelve σ parameters to calibrate, describing transition rates for disease progression, with six for each age. These define the probability of moving to the next worse health status and the probability of recovery at the three infectious stages: asymptomatic, feverish, and hospitalized. Our calibration will impose that the old are 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, we identify the six values for σ from the following six target moments: the average length 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. We set the duration of the asymptomatic phase to 5.1 days and

¹⁹We categorize food away from home, transportation services, apparel, new vehicles, and gasoline as luxuries, and the remaining categories as basic goods. We treat medical care as part of the basic sector, even though it involves social contact, because shutdowns have largely exempted healthcare.

²⁰Total spending in 2019 was \$7,094bn, of which Social Security spending and Medicare spending were \$1,031bn and \$786bn respectively (NIPA Tables 3.1 and 3.12).

assume this is common across age groups.²¹ We assume a common-across-age-groups average duration of the fever state of 7 days. Following current CDC estimates, we allow the average duration of the hospitalized state to vary by age, with average duration of 6.2 days for the young and 8.1 days for the old.²²

The exit rate from the asymptomatic state to recovery defines the number of asymptomatic cases of COVID-19. This parameter is important because if a large share of infected individuals never develop symptoms, the overall infection-fatality rate for the virus will be low. Following Buitrago-Garcia et al. (2020) we assume that 31 percent of those infected recover without ever developing symptoms.²³ This implies that roughly half of those actively infected at a point in time will be symptomatic.

We set age-specific death probabilities, conditional on being infected, to target (i) an overall (population-weighted) infection fatality ratio (IFR) of 1.0 percent as of March 21, 2020, and (ii) an IFR for the old that is 40 times that for the young. With regard to the first target, a 1.0 percent IFR is consistent with existing estimates. For example, Brazeau et al. (2020) estimate an overall IFR of 1.15 percent for high income countries. In addition, we will verify that this IFR delivers reasonable model predictions for cumulative infections through 2020 (conditional on matching cumulative deaths). The age differential in the IFR reflects several pieces of evidence. First, on a per capita basis COVID has killed 21 times more people over age 65 than people below that age. Second, even that differential understates the true impact of age on COVID lethality, since seroprevalence studies indicate that those over 65 have been infected at only around half the rate of younger individuals (see, for example, (Bajema et al., 2021)), a fact also true in our model. A factor of 40 differential is consistent with the international evidence presented in O’Driscoll et al. (2021). We will also show that in the context of our model, it delivers a realistic share of deaths accruing to the old versus the young. Thus, we target an IFR for the young of $0.01/(0.85 + 0.15 \times 40) = 0.146$ percent and an IFR for the old of $40 \times 0.146 = 5.84$ percent.

Within our model the overall age-specific IFR is the product of the probability of becoming

²¹Note that 5.1 days is strictly an estimate of the incubation period for COVID (Lauer et al., 2020), and for the first few days of the incubation period, people who have been exposed are likely not contagious. Our model does not differentiate between an “exposed but not contagious” phase and a “pre-symptomatic but contagious” phase. Still, the baseline (CDC, 2020, Table 1) estimate is that 50 percent of COVID transmission occurs before symptom onset.

²²See Table 2 of CDC (2020).

²³The CDC (2020, Table 1) estimates that 40 percent of infections are asymptomatic, but that asymptomatic individuals are only 75 percent as infectious as symptomatic individuals.

symptomatic, 0.69, times the age-specific probability that fever becomes so severe to require hospitalization, times the age-specific probability of death conditional on hospitalization. Of the latter two probabilities, the probabilities of death conditional on COVID hospitalization are the best measured. The CDC (CDC, 2020, Table 2) reports a death probability of 6.2 percent for younger hospitalized individuals and of 26.6 percent for those over 65.²⁴ Given these death rates and our IFR targets, the implied age-specific probabilities of being hospitalized are 3.41 percent for the young and 31.8 percent for the old.

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 through April 12. This implies $\beta_e = 0.80$.²⁵ The values of α_w , α_c , and β_h determine the overall basic reproduction number R_0 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 percent of potentially infectious inter-person contact happens in workplaces and schools, 19 percent occurs in travel and leisure activities, and the remainder takes place at home and in other settings. 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 . Note that this evidence does not pin down the *levels* of α_w , α_c , and β_h , or equivalently the level of R_0 . We set the level for these infection-generating-rate parameters to deliver an *initial* R_0 of 2.5, which was the CDC's best estimate for COVID-19 at the time of writing (CDC, 2020, Table 1).²⁶

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 percent of COVID-19 hospital admissions require intensive care, suggesting that total hospital capacity is around

²⁴The CDC reports death rates of 2.4 percent and 10.0 percent for the 18 – 49 and 50 – 64 age groups. Our 6.2 percent rate is a simple average, which is consistent with the relative shares of those groups hospitalized. See Covid-Net at https://gis.cdc.gov/grasp/covidnet/covid19_5.html.

²⁵5% is an estimate by Sepkowitz (2020) of the share of infections accruing to health-care workers who acquired the infection after occupational exposure. As of March 24th 2020, 14% of Spain's confirmed cases were health care workers (New York Times, 2020).

²⁶See the Appendix for more details.

58,000/0.215=270,000. Tsai et al. (2020) emphasize that this capacity is very unevenly allocated geographically, 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 0.03 percent of the population is hospitalized. We set the parameter λ_o so that the average mortality rate in emergency care when hospitalizations reach 200,000 is 25 percent above its value when capacity is not exceeded.²⁷

[Table 1 About Here]

4.2 Time Varying Parameters

Baseline Mitigation Our baseline path for mitigation, designed to approximate historical U.S. policy, assumes no mitigation ($m_t = 0$) until March 20, with shutdowns starting on March 21. California announced the closure of non-essential businesses on March 19, and New York and Illinois did so on March 20. From March 21 the path for m_t is chosen so that the model replicates the dynamics of employment from March through December 2020.²⁸

Improved Treatment Treatment for hospitalized COVID patients improved during the course of 2020. In part, this reflected the introduction of new therapies, such as the steroid Dexamethasone and the antiviral Remdesivir. Perhaps more important was steady refinement of best practices related, for example, to optimal use and calibration of mechanical ventilators. The Institute for Health Metrics and Evaluation (2020) estimates that hospital mortality declined by 30 percent between March and September. We assume that mortality rates decline steadily and geometrically, beginning on March 21, so that $\sigma_{t+k}^{ed} = \exp(-\delta k)\sigma_t^{ed}$ with $\delta = 0.71$ on an annual basis. Thus the overall IFR falls from 1 percent on March 21 2020 to $\exp(-0.71 \times 0.78) = 0.57$ percent by December 31, 2020.

Time Variation in R_0 It is well understood that the basic reproduction rate is not a biological constant but rather something that varies with environmental conditions and behavior.

²⁷Much of the early concern about exceeding capacity 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.

²⁸We target the series for the Employment-Population Ratio [EMRATIO] published by the U.S. Bureau of Labor Statistics, where we index the series to equal one in February 2020. We estimate a logistic form for mitigation to fit this series, which yields a baseline mitigation function of $m(t) = 0.653/(1 + e^{0.008(t-50)})$, where t is days since March 21. Figure D.3 in the Appendix plots the implied model employment rate against the BLS series.

We will model this variation in a simple way, allowing for two sources of changes in R_0 . We implement this variation via a shifter ζ_t that proportionately scales all the infection-generating-rate parameters α_w , α_c , β_h , and β_e (and thus ζ_t also scales R_0). Before March 21 2020, we normalize $\zeta_t = 1$.

The first source of variation in R_0 is that we assume that people changed their behavior in a permanent fashion on March 21, the same date we initiate shutdowns. At this date, more cautious behavior leads to a decline of ζ_t to $\zeta_H < 1$.

Second, there is evidence of strong seasonality in COVID transmission, with much faster transmission in colder months than warmer ones. For example, the time paths for COVID deaths in Europe versus South America are mirror images of each other. This may reflect the impact of temperature or humidity on virus spread. Or it may simply reflect the fact that when temperatures are cold, people spend more time indoors, where ventilation is worse and transmission is therefore easier. We assume infection-generating rates fluctuate seasonally, being relatively high for half the year and relatively low during the other half. In the colder, high-transmission season, $\zeta_t = \zeta_H > \zeta_L$. In the warmer, low-transmission season, $\zeta_t = \zeta_L$.

This modeling strategy is summarized by four parameters: (i) ζ_H , (ii) ζ_L , (iii) the six-month-apart dates when the seasons change, and (iv) the number of seed infections at the date the simulation starts. We set these four parameters to match cumulative official deaths at four dates: April 12, May 31, October 31, and December 31, 2020, which were 27,003, 108,102, 238,347 and 371,503, respectively.²⁹ We focus on matching deaths, since deaths are the best-measured and most welfare-relevant measure of the virus' impact.³⁰ The logic for these choices of dates is that April 12, Easter, was the peak of the first wave of deaths. Daily deaths were then relatively stable and relatively low from the end of May to the end of October, when the Fall surge began in earnest.³¹ We start our simulation on February 1, 2020. To generate 27,003 deaths by April 12, given that start date and an initial R_0 of 2.5, requires 647 initial infections

²⁹All our data numbers for deaths are from the Centers for Disease Control: <https://covid.cdc.gov/covid-data-tracker/#datatracker-home>.

³⁰Following most of the existing literature, we focus on replicating official deaths, while recognizing that if official statistics understate true COVID mortality, the case for mitigation would be strengthened. Official COVID deaths were 345,700 up through the end of December, but excess death measures suggest this official tally may fall up to 100,000 short of the true toll. See, for example, Rossen et al. (2020). These excess deaths are very highly correlated temporally and geographically with official COVID deaths (Woolf et al., 2020).

³¹There was also a smaller mid-summer surge, concentrated in sunbelt states. This may have reflected high temperatures in those states, driving people to seek air-conditioning indoors.

on February 1, 2020.

This calibration strategy yields a decline in R_0 from 2.5 before March 21 to 1.26 after March 21, reflecting a value for ζ_H of 0.56.³² The effective R_t (including the effects of mitigation) declines from 1.26 to 0.96 on March 21 when mitigation starts. Thereafter the dynamics of R_t are influenced by two countervailing trends: gradually declining mitigation pushes R_t up, while gradual growth in the size of the recovered population pushes R_t down. On top of that is the impact of seasonality. When ζ falls to $\zeta_L = 0.47$ on April 22 transmission slows, with R_0 switching to 1.13. This seasonality allows the model to replicate the observed relatively modest increase in cumulative deaths between June 1 2020 and the end of October, as well as the observed surge in mortality in November and December.

[Table 2 about here]

We assume that $\bar{T} = 470$ days after the start of the pandemic, the economy is back in steady state: all those who survived the pandemic are healthy, and all young survivors are working.³³ All epidemiological and economic parameter values are summarized in Tables 1 and 2. The calibration implies the Spring 2020 population health distribution described in Table C.1 of Appendix C.2. On March 21, 0.82 percent of the U.S. population was actively infected (including asymptomatic infections), with that number rising to 0.84 percent by April 12, with an additional 1.82 percent having recovered.³⁴

5 Optimal Mitigation During the First Phase of the Pandemic: 2020

We now describe the quantitative findings from our model. We start in Section 5.1 by describing the model's predictions for health outcomes under the benchmark mitigation policy. The assumed benchmark fiscal constitution includes access to government debt and distortionary labor income taxation. We assume the paths for taxes and transfers to non-workers are chosen optimally (as characterized in Section 3) starting at the time the COVID-19 pandemic begins

³²Recall that all the R_0 values we report in the text are the values that would obtain absent any economic mitigation (and, in the standard way, the ones that obtain when the entire population is susceptible). In Figure D.1 in the Appendix we plot the time path for a different concept of R_0 , which incorporates the effects of economic mitigation (but still assumes a fully susceptible population).

³³Thus, the pandemic ends on July 4, 2021, the day President Biden celebrated "independence from the Coronavirus." At the onset of the COVID-19 pandemic there was of course considerable uncertainty about how long it would last.

³⁴These numbers are within the range of expert estimates from the COVID-19 survey compiled by McAndrew (2020) at the University of Massachusetts.

(March 21, 2020).³⁵ We explore the model’s economic predictions in Section 5.2. Section 5.3 characterizes the optimal mitigation path chosen by a utilitarian government. Section 5.4 shows that the utilitarian optimal policy masks large differences in policy preference across different household types. Section 5.5 turns to the interaction between mitigation and fiscal policy. We stress that very different mitigation paths are optimal depending on the fiscal regime, and specifically, whether the government has access to government debt and lump-sum taxes. We argue that representative-agent models implicitly assume costless redistribution, and that those models therefore exaggerate optimal mitigation when redistribution is costly. Section 5.6 explores how age variation in the health threat posed by COVID impacts the optimal path for mitigation by contrasting our baseline model to a “representative health” alternative. After having validated the model against data from 2020, in Section 6 we characterize the optimal mitigation policy from January 1 2021 onward, as vaccines are deployed.

5.1 Health Outcomes Through the Lens of the Model

In Figure 1 we display U.S. times series for daily deaths and currently hospitalized individuals alongside their model counterparts.³⁶ Observe that the model tracks actual COVID-19-related deaths very well, both in the spring of 2020 and in the fall/winter 2020 wave. The model also quite closely replicates the path for hospitalizations, but overstates hospitalizations in the fall of 2020. This is likely due to the fact that in the model *all* individuals that experience severe COVID-19 and eventually succumb to the virus have to move through the emergency hospital E state, whereas in the data some very sick people were not hospitalized but died in nursing homes and other settings.

[Figure 1 about here]

By July 4th, 2021 the model predicts 660,000 cumulative deaths (see Table 3), corresponding to 0.2 percent of the U.S population. For comparison, the Institute for Health Metrics and Evaluation (IHME) reports 605,000 actual deaths as of that date. Of these deaths, 66,000 reflect

³⁵Quantitatively, it does not matter much whether the government optimally adjusts the level of taxes and transfers at the beginning of the pandemic, or keeps taxes and transfers constant at their pre-pandemic levels throughout the pandemic and then optimally adjusts those variables when the pandemic ends, in period \bar{T} . Appendix D.2 contains results for this alternative specification.

³⁶Both data series are from the Centers for Disease Control. The series for hospitalizations was accessed from <https://covid.cdc.gov/covid-data-tracker/#hospitalizations> on June 15, 2021 and updated as new data was released. The data for daily deaths is at https://covid.cdc.gov/covid-data-tracker/#trends_dailytrendscases.

hospital overload.³⁷

One of our central themes is that these health outcomes were very different for the young versus the old. This is demonstrated in the top four panels of Figure 2 which displays the model health dynamics for the three groups: basic workers, luxury workers, and the old. The top left panel shows the currently infected percentage of the population and the top right panel plots the currently hospitalized share of each group. The left panel in the second row depicts the total share of the population groups that has had COVID-19 and recovered, whereas the right panel in the second row shows the population shares that have died from the disease.

[Figure 2 about here]

One can clearly see how the young, especially those in the basic sector (who are never mitigated and care for the hospitalized), are more likely to be infected than the old (who do not risk workplace infection). But a much larger share of the elderly population is hospitalized (around 0.25 percent of the population 65 and older at the peak of the winter 2020 wave). This differential hospitalization rate contributes to a massive gap in the incidence of death from COVID-19: according to our model, by July 4, 2021 more than 1 percent of the elderly population has died from COVID, whereas among the young that number is only approximately 0.035 percent. Through the end of 2020, 81 percent of model COVID-19 deaths are accounted for by the old. This accords well with the data for the U.S., according to which 80 percent of COVID-19 deaths in 2020 were among people aged 65 and older.³⁸

By December 31, 2020, around 30 percent of basic workers, 25 percent of luxury workers, and 14 percent of the elderly have recovered, and in the aggregate, slightly less than 24 percent of the model population (79.1 million people) has had a COVID-19 infection. Although there is a wide range of estimates for the corresponding number in the data, the prediction of our model is roughly in the middle of the estimates for the United States of which we are aware.³⁹

³⁷We compute this by running a simulation in which the overload parameter λ_o in equation 12 is set to zero.

³⁸See CDC weekly updates: https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm.

³⁹The Institute for Health Metrics and Evaluation (IHME) estimates 18% percent of the U.S. had been infected as of December 20, and the AI-based model by Youyang Gu (<https://covid19-projections.com/>) gives a point estimate of 23.6%, with a range of 15.7% to 35.4% for December 31. On the high end, the CDC estimates that 28% of the total population had COVID-19 by December 1, 2020 (<https://twitter.com/youyanggu/status/1344002411556339712>).

5.2 The Economic Crisis Through the Lens of the Model

We now turn to the economic dimension of the pandemic. The bottom four panels of Figure 2 characterizes the economic recession, under the baseline shutdown policy. It plots GDP, the optimally chosen tax rate and per capita transfers, and government debt (the latter two normalized by pre-COVID GDP).

Under our estimated empirical mitigation path, the government shuts down 40 percent of the luxury sector. This plunges the economy into a deep recession in March 2020 (third row, left panel). The impact of mitigation is amplified by a rise in the share of people off work because they are sick or hospitalized. The reduction in economic activity reduces tax revenue at the same time the government must finance transfers to a larger share of the population. But the government has access to debt and understands that this extra fiscal stress is temporary. Thus, the optimal fiscal response is reduce transfers and increase the tax rate only very mildly – albeit permanently – and to finance most of the extra spending by issuing debt. These optimal fiscal responses are characterized theoretically in Section 3.1. The permanent increase in the income tax rate is only 0.21 percentage points, and thus the tax-induced reduction in labor supply for those working is modest. As Section 5.5 will demonstrate, this is in stark contrast to an economy in which the government does not have the fiscal space to run a deficit and therefore needs to raise taxes or lower transfers significantly during the pandemic recession. That in turn will curb the appetite of the government to engage in economically costly mitigation efforts to contain the pandemic.

As time passes and mitigation is gradually relaxed, output slowly recovers. However, when the fall/winter wave hits, the recovery slows down in late 2020 and early 2021, reflecting the fact that more people are sick and not working.⁴⁰ The total increase in debt in our simulation during the first 12 months of the pandemic amounts to about 3.5 percent of pre-COVID GDP. Because government deficits are financed by borrowing from abroad, the average trade deficit in the model (defined as $Y_t - C_t - G_t$) is similar. For comparison the total increase in spending on unemployment insurance in 2020 and 2021 (relative to a 2019 baseline) amounted to 3.75 percent of 2019 U.S. GDP, while the U.S. trade deficit widened by 0.8 percent of 2019 GDP.⁴¹

⁴⁰The discrete jump in output on 7/4/2021 reflects our assumption in this section that the pandemic ends on that date.

⁴¹Unemployment insurance spending is from NIPA Table 3.12, line 7. The U.S. trade balance on goods and services is from FRED series BOPGSTB and GDP. The balance from April 2020 through March 2021 was \$177bn larger than in the previous 12 month period.

5.3 Consequences of Alternative Mitigation Paths

What mitigation policy is optimal? In Figure 3 we characterize the optimal policy for a utilitarian government, and contrast the implied outcomes to those under the baseline mitigation path. Overall, the extent of shutdowns under the optimal policy is similar to that under the baseline path, but there are significant differences in timing.⁴² In particular, shutdowns are optimally slightly milder early in the pandemic, but the optimal mitigation path inherits some of the seasonality in the infection-generating rate, with shutdowns ramping up in the fall.

[Figure 3 about here]

Since optimal mitigation is initially lower than under the actual policy, the paths for hospitalizations and deaths at first run higher. But by reimposing moderate shutdowns in the fall, the optimal policy avoids the surge in hospitalizations and deaths that occurs under the baseline (and in the data). Because the optimal path for mitigation is flatter relative to the baseline, so is the path for GDP. The time series for transfers and tax rates are similar in the two economies, though the fact that mitigation is more prolonged under the optimal policy translates into a slightly higher optimal tax rate and lower transfers.

Table 3 reports various outcomes through July 4, 2021 under three different mitigation paths: the baseline, meant to capture actual policy, the utilitarian-optimal path, and a policy of zero mitigation. Relative to the baseline, had the U.S. instead followed the optimal mitigation path, the model predicts we would suffered more deaths in 2020 (Figure 3), suggesting lockdowns were initially too tight, but 300,000 *fewer* deaths in total, as tighter restrictions in late 2020 and early 2021 would have largely eliminated the winter wave.⁴³

[Table 3 about here]

⁴²This implies that the Pareto weights that would deliver the *initial* empirical size of lockdowns as optimal policy are close to the utilitarian welfare weights we actually use. Section 5.5 will demonstrate that the availability of government debt is important for this finding.

⁴³Recall that sub-sectors in our model economy vary in terms of infection-generating contact rates. Our baseline assumption is that mitigation is directed toward the riskiest sub-sectors. We have also explored an alternative model for mitigation according to which sub-sectors are shut down at random. In this undirected mitigation model, the terms $\alpha_w(1 - m_t)$ and $\alpha_c(1 - m_t)$ in equations (7) and (8) are replaced by α_w and α_c . This implies higher infection rates for any interior level of mitigation, but does not change the economic cost of mitigation. We find that with undirected mitigation, the utilitarian optimal path for mitigation m_t is always higher than in the baseline directed-mitigation model. It is perhaps surprising that the government chooses more mitigation when mitigation is a less effective tool. The reason is that when mitigation is less effective, infections and hospitalizations run higher. Absent higher mitigation, hospital capacity would be drastically exceeded, leading to even higher mortality.

A policy of no mitigation at all would have led to much higher mortality. Interestingly, that is not because there would have been more cumulative infections: in fact cumulative infections are very similar to those under the baseline policy. Instead, a policy of no mitigation leads to a huge initial surge of infections and hospitalizations. Mortality is much higher than in the baseline because (i) hospitals are overloaded, and (ii) infections occur earlier in time, before treatment has improved. On the positive side, a policy of no mitigation would have meant a much smaller recession, and thus a lower economic cost of the pandemic.

5.4 Preferences over Optimal Mitigation: Heterogeneity by Age and Sector

We now explore how preferences for mitigation vary by age and by sector. Figure D.2 in Appendix D.1 describes the mitigation paths that each type of household would choose at the start of the pandemic, if they were allowed to dictate policy. The main message is that the young and the old have sharply diverging policy preferences. The old would prefer a much more extensive and long-lasting lockdown, initially covering 100 percent of the luxury sector and lasting into the spring of 2021. In contrast, the young would like to see much milder shutdowns, and they would like to see lockdowns end altogether before the end of 2020. The policy chosen by the utilitarian government is a compromise between these divergent preferences.⁴⁴ Comparing the policies preferred by young workers in the basic and the luxury sectors, basic workers favor slightly more extensive lockdowns, because they do not risk losing their jobs from mitigation and because they are at higher risk of getting sick.

Table 4 reports welfare gains for different individual types (the rows) under different mitigation paths (the columns), relative to the empirical mitigation path. The welfare calculations ask: what percent of consumption would a person be willing to pay every day for the rest of her life to move from the baseline mitigation path to either zero mitigation or to one of the four mitigation paths plotted in Figure D.2?⁴⁵

⁴⁴There are two reasons to mitigate in our economy. First, mitigation saves lives. Second, with fewer infections, fewer people feel miserable with flu-like symptoms or suffer in hospital, which enters welfare via the preference shifters \hat{u}^f and \hat{u}^e . We have also solved for the utilitarian-optimal policy for a calibration in which we set these preference shifters to zero. We found utilitarian optimal mitigation in this case to be very similar to the baseline specification, indicating that the rationale for lockdowns is almost entirely to save lives. The logic is that even though many more people suffer COVID symptoms than die from the disease, symptoms are very short-lived relative to the duration of flow utility lost from a premature COVID death. Thus, if a hospital treatment for COVID had emerged that would completely eliminate mortality, it would be optimal to set mitigation to zero, according to our model.

⁴⁵Let stars denote allocations under an alternative policy. The welfare gain for the young from switching to the

[Table 4 about here]

The table reinforces a key message of the paper: the young and the old have sharply divergent policy preferences. The reasons are clear. The old worry about health, but not about higher taxes or job loss, and thus they favor strict lockdowns. The young worry less about health and much more about the higher taxes and job losses that mitigation entails.

In addition, the magnitudes of welfare gains and losses across alternative policies vary strongly by age, with much more at stake for the old than for the young. The old face disproportionate risk of severe illness and death in a virus surge, which translates to losses exceeding 6 percent of permanent consumption in the counter-factual scenario with zero mitigation. Conversely, the old see welfare gains approaching 7 percent under their preferred extensive and persistent mitigation policy.

In contrast, the welfare effects of alternative policies are modest for basic sector workers. They never worry much about health, and do not risk losing their jobs in a shutdown. Luxury sector workers suffer more from the extensive mitigation favored by the old, since they bear all risk of job loss. Overall, all young individuals would modestly prefer zero mitigation to the baseline path. Conversely, the larger and more persistent lockdown preferred by the old would substantially reduce their welfare.⁴⁶

alternative policy from the baseline policy is the value for ω that solves

$$\begin{aligned} & E_0 \left\{ \int_0^{T^y} e^{-\rho t} S_t^y \left[\log \left((1 + \omega) \left(c_t - \frac{\varphi}{1 + \frac{1}{\chi}} h_t^{1 + \frac{1}{\chi}} \right) \right) + \bar{u} + \widehat{u}_t^i \right] dt \right\} \\ &= E_0^* \left\{ \int_0^{T^y} e^{-\rho t} S_t^{y*} \left[\log \left(c_t^* - \frac{\varphi}{1 + \frac{1}{\chi}} h_t^{*(1 + \frac{1}{\chi})} \right) + \bar{u} + \widehat{u}_t^{i*} \right] dt \right\}. \end{aligned}$$

Note that the equilibrium allocations in our economy (as is standard with GHH preferences) have the property that

$$c_t - \frac{\varphi}{1 + \frac{1}{\chi}} h_t^{1 + \frac{1}{\chi}} = \frac{c_t}{1 + \chi}.$$

Thus the welfare gain ω can be interpreted as the percentage change in consumption that leaves a young individual indifferent between the two policies. The same argument applies to old individuals.

⁴⁶The strong age heterogeneity in the preferences over mitigation raises the question of whether age-based policies could yield welfare gains; see Acemoglu et al. 2021 and Brotherhood et al. 2020 for an extensive discussion of such policies. In a previous version of this paper, Glover et al. 2020a we considered one such policy: a policy to eliminate elderly infections while shopping by asking the young to shop on their behalf. We found significant welfare gains for the old.

5.5 The Interaction between Mitigation and Redistribution

The second main theme of this paper is that mitigation and redistribution policies interact. Shutdowns increase unemployment and thus necessitate more redistribution. The set of fiscal tools available to the government will determine the cost of this redistribution, and will therefore influence the government's appetite for mitigation.

To quantify the importance of this point we now compare optimal mitigation policies under the benchmark fiscal constitution against two alternatives emerging from either permitting the government access to lump-sum taxation or removing its access to government debt. An economy with lump-sum taxation (with or without debt) is an attractive point of comparison because with lump-sum taxes, flow utilities of all groups are equalized under the optimal tax-transfer policy.⁴⁷ Because there are no distributional conflicts, the economic side of the model is effectively a representative household economy. The constitution in which the government does not have access to debt is informative about how countries should optimally adapt policy if they lack the fiscal space to smooth the economic costs of shutdowns over time, an issue also explored by Arellano et al. (2023).

The top left panel of Figure 4 shows the utilitarian-optimal mitigation paths for three fiscal constitutions: (1) the baseline with distortionary taxes and government debt (solid blue line), (2) the lump-sum tax economy with debt (dashed orange), and (3) distortionary taxation but no access to debt (dotted red). We make four broad observations. First, the government chooses more mitigation when it can borrow than when it cannot. Second, optimal mitigation is higher when the government has access to lump-sum taxes than when it does not. Third, the impact of access to debt is quantitatively more important than the availability of lump-sum taxation. Fourth, the optimal mitigation path inherits the seasonality of the health pandemic under all fiscal constitutions, but the seasonality is less pronounced when the government can borrow.

As in Lucas and Stokey (1983), access to debt permits the government to smooth taxes and transfers over time, and to separate the timing of mitigation (which is optimally time-varying) from the timing of taxes and transfers, both of which are optimally smooth; see Proposition 1 and the middle two panels of Figure 4. A small but persistent increase in taxes is much less

⁴⁷Without GHH preferences, the government would only equalize marginal utilities. Given our GHH specification, both marginal utilities and utility *levels* are equalized. If the government had access to both lump sum taxes/transfers and distortionary labor income taxes, it would not use the distortionary labor income tax, as is always the case in representative agent economies.

costly than a large temporary increase, which incentivizes the government with access to debt to pursue more extensive shutdowns.⁴⁸ Of course, the flip side is that the government accumulates significant extra debt (close to 4 percent of pre-COVID GDP) which has to be repaid after the pandemic is over. Consequently taxes are higher and transfers are lower in baseline economy in the long run, relative to the corresponding no-debt (PAYGO) economy. However, as the figure demonstrates, the magnitude of this effect on the tax rate is quantitatively very small, since the COVID-19 crisis is short relative to the horizon over which the higher debt has to be repaid): the tax rate has to rise by only 0.2 percentage points, from 30.4 percent to 30.6 percent.

[Figure 4 about here]

Optimal mitigation is lower in the baseline distortionary tax economy relative to the alternative lump-sum tax specification in which there are no inequality concerns or distortions. This was the content of Proposition 3.⁴⁹ The logic is that with distortionary taxes, higher tax rates translate into larger distortions to labor supply, and the government faces an equity-efficiency trade-off in setting taxes and transfers. In optimally balancing this trade-off, the government sets transfers to a level such that workers enjoy higher utility than non-workers. By creating more non-workers, mitigation thus directly increases inequality. In addition, since more mitigation requires higher tax rates, implying a higher efficiency cost of taxation, the government also optimally reduces transfers, further widening inequality in effective consumption between workers and non-workers. With lump-sum taxes instead the government can increase mitigation without lowering productive efficiency and without amplifying inequality.

Note that although the government mitigates less with distortionary taxes, debt rises by more as a share of GDP. That is because transfers are higher with distortionary taxes while output is lower, so deficits relative to output are more sensitive to the level of mitigation.⁵⁰

⁴⁸One reason smoothing taxes is desirable is that the marginal distortion associated with a higher proportional tax rate is increasing in the tax rate. Another reason is that households in our model have no private access to credit. Thus, smoothing taxes helps households smooth consumption over time.

⁴⁹Note that the proposition does not imply that optimal mitigation should be *uniformly* lower in the distortionary tax economy. The reason is that different paths for mitigation imply different health state distribution dynamics in the two economies and thus differential marginal gains from additional mitigation in terms of improved health outcomes.

⁵⁰The fact that GDP is lower with distortionary taxes than with lump-sum taxes reflects the fact that hours worked are a declining function of the after-tax wage. Two of our modeling assumptions play a role in explaining why transfers are *higher* with distortionary taxes. First, we assume government purchases are a fixed fraction g of output. Because output is higher with lump-sum taxes, total government spending is higher in that economy, and utility levels therefore tend to be lower. For non-workers, lower utility means lower transfers. The second important assumption is our non-separable GHH utility function. Because workers work longer hours with lump-

Taken together these comparisons clearly illustrate that mitigation and fiscal flexibility are complementary. When the government has more fiscal tools, stronger mitigation is optimal. There is value to being able to redistribute resources costlessly over time (via access to debt) and to being able to costlessly redistribute resources intra-temporally from workers to non-workers (via lump-sum taxes). Switching on or off either instrument changes the optimal path for mitigation, which in turn has a significant impact on the dynamics for COVID mortality; see again the bottom left panel of Figure 4.

5.6 Why Does Household Heterogeneity Matter for Optimal Mitigation?

Proposition 3 answered this question theoretically, and the last section quantified it for the economic dimension of household inequality. This section argues that explicitly modeling heterogeneity in the health risks posed by the pandemic, especially heterogeneity by age, is equally crucial for the optimal mitigation response. We compare our baseline economy to a counterfactual alternative in which young and old agents face the same risk of being hospitalized following a COVID infection, and, conditional on being hospitalized, face the same risk of death. We call this counterfactual the “representative health” economy. We set the probabilities of being hospitalized and of dying in the representative health economy to the age-group-weighted averages of the respective age-specific probabilities reported in Table 1. Thus, the probability any symptomatic individual requires hospitalization in the representative health economy is $0.85 \times 3.41 + 0.15 \times 31.8 = 7.7$ percent, while the probability a hospitalized individual dies at the start of the pandemic is $0.85 \times 6.2 + 0.15 \times 26.6 = 9.3$ percent. All other parameters in the representative health economy are identical to those in the baseline.

[Figure 5 about here]

Figure 5 contrasts outcomes in the representative health economy to those in the baseline. The top left panel shows daily deaths in the two economies, given the same empirical path for mitigation in both economies. Deaths run somewhat lower in the representative health economy. The reason is that, while the two economies deliver similar hospitalization rates, in the baseline economy the people most likely to be hospitalized (the old) are also the most likely to die conditional on hospitalization, while this positive correlation is broken in the representative health economy.

sum taxes, the government wants to give them higher consumption to equate utilities across individuals in the lump-sum economy, translating into lower transfers to non-workers.

The top right panel of the figure reports, on the same daily basis, total years of life lost due to COVID deaths. Here we see the opposite picture: the total number of years of life lost due to COVID is larger in the representative health model, even though COVID kills fewer people. The reason is simple: in the baseline model, deaths are heavily concentrated among the old who have relatively short remaining expected lifetimes (14 years, relative to 48 years for the young), whereas in the representative health model, the young and old have similar expected mortality.⁵¹

The middle left panel compares the utilitarian optimal mitigation paths in the representative health and the baseline economies. The government chooses significantly higher mitigation in the representative health economy. The main reason is that the government internalizes that individuals enjoy a flow value from being alive, and the government's incentive to mitigate is therefore tied more closely to years of life lost to COVID than to the raw mortality rate. A second reason optimal mitigation is higher in the representative health model is that mitigation directly reduces workplace infections, which are infections of young people. Reducing infections among young people is more worthwhile when the young face more adverse health risks from COVID. The middle right panel of the figure shows that higher mitigation comes at the cost of a larger decline in economic activity.

The bottom two panels of Figure 5 report daily deaths and daily years of life lost due to COVID in each economy, under the respectively optimal paths for mitigation. The harsher lockdowns that are optimal in the representative health economy translate to a much lower death toll from COVID. But the total life years lost from COVID are quite similar in the two economies.

The lesson we take from the last two subsections is that heterogeneity matters for the optimal policy response to COVID-19. Policy recommendations based on representative-agent models that abstract from the connection between shutdowns and economic inequality and from heterogeneity in the health impact of COVID will tend to prescribe excessively severe shutdowns, for two reasons. First, such models abstract from the fact that health risks are concentrated on the old. Second, they miss that shutdowns exacerbate economic inequality and amplify the distortions associated with redistribution.

⁵¹The young are actually slightly more likely to die from COVID than the old, because they are infected at a higher rate (recall the old do not work), but this difference is quantitatively minor.

6 Optimal Policy in the Presence of a Vaccine Roll-Out

By the beginning of 2021, it was clear that a roll-out of effective vaccines was imminent. We now study optimal mitigation and redistribution policy when a vaccine becomes available. When susceptible people are vaccinated, we assume they immediately transition to the recovered state in the epidemiological model, thus pushing the economy closer to herd immunity.⁵² Vaccines are rolled out among the different model age groups in accordance with the actual pattern for the U.S. in the first half of 2021.⁵³

We make one further change for our 2021 simulations relative to those for 2020, which is to assume that as vaccinations proceed, the infection-generating scaling parameter ζ_t gradually rises in proportion to the share of newly recovered individuals. The idea is that as the vaccination rate rises, people become less concerned about contracting and transmitting the virus and return to pre-COVID social behavior. We set the constant of proportionality so that if everyone is vaccinated – implying a fully recovered population – the basic reproduction number R_0 would return to its initial, pre-March 20, 2020 value.⁵⁴

Our simulations in this section start on January 1, 2021. The initial condition is the population health distribution for that date implied by the simulation described in Section 5 under the benchmark mitigation path, which we argued to be a good approximation to the actual distribution at the end of 2020. We also take the initial value for debt from that simulation.

The top-left panel of Figure 6 displays mitigation paths for 2021 preferred by a utilitarian government in the presence of a vaccine roll-out (red dotted line) and in the absence of one (solid blue line).⁵⁵ The main takeaway is that the optimal level of mitigation in early 2021 is initially significantly higher in the presence of the vaccine relative to the no-vaccine scenario. But reopening occurs earlier with a vaccine: the economy is fully open by mid-February 2021, whereas

⁵²Recent evidence suggests that fully vaccinated individuals can still transmit the virus, especially the “Delta” and “Omicron” variants, though they are less likely to do so than infected but unvaccinated individuals. We explore vaccine deployment in the context of the Delta variant in Glover et al. (2022), where we study the implications of alternative vaccination strategies that strictly prioritize either the old or the young in vaccine distribution.

⁵³Model individuals (including those who have already recovered from COVID) are vaccinated over the six month period between January 1 and July 1, 2021. We assume that 0.47 percent of the old and 0.30 percent of the young are vaccinated per day, which approximately replicates the differential pace of vaccination rates by age reported by the CDC (see <https://covid.cdc.gov/covid-data-tracker/vaccination-demographics-trends>). We assume that vaccinations of the young and old continue at a constant pace until 60 percent of the young and 80 percent of the old have recovered, either by surviving infection or being vaccinated.

⁵⁴See Figure D.1 in the Appendix for the model-implied evolution of the basic reproduction number.

⁵⁵We impose the same path for ζ_t in the no vaccine scenario as in the vaccine scenario.

significant mitigation optimally persists until the end of May absent a vaccine. The logic for a harder early lockdown when vaccines are known to be coming is that infections prevented by mitigation in early 2021 will never occur in the future. Absent a vaccine, in contrast, mitigation primarily delays infections rather than eliminating them altogether.⁵⁶

[Figure 6 about here]

How the availability of vaccines changes the trajectory of economic variables is illustrated in the upper-right and middle two panels of Figure 6, which show the evolution of GDP (relative to its pre-COVID level), the optimal tax rate, and the implied debt-to-GDP ratio. In both scenarios, GDP is initially depressed relative to pre-COVID times because part of the economy is shuttered, because some workers are out sick, and because a higher tax rate reduces labor supply. Of these three effects, the direct effect of shutdowns is quantitatively the most important, and the paths for GDP therefore closely track the inverse of the chosen mitigation paths.

In the scenario with vaccines, optimal mitigation implies a temporary collapse in economic activity, followed by a swift recovery, as shutdowns are ended. In contrast, absent a vaccine, there is a smaller initial decline in GDP but a more prolonged recession into the spring of 2021. Thus, from the perspective of our model, getting the pandemic under control with a vaccine is the key to a rapid economic recovery. Debt dynamics are also largely driven by the chosen paths for mitigation.

The bottom two panels of Figure 6 show daily deaths with and without vaccines, and the share of recovered individuals, where transition into the recovered state can either occur through vaccination or by contracting and surviving the disease. Daily deaths are significantly lower during the spring of 2021 with a vaccine, and again lower towards the end of the summer, when herd immunity has been reached with vaccines, but not in the no-vaccine scenario.

6.1 Welfare Gains from Vaccination

Table 5 summarizes the welfare gains from the introduction of a vaccine, always relative to a world where no vaccine ever emerges and mitigation is set to the empirical benchmark path. The different columns of the table correspond to three different scenarios. In the first scenario,

⁵⁶Note that relative to the spring and summer of 2020, optimal utilitarian mitigation is significantly higher in early 2021 even in the absence of a vaccine. This is due to the fact that the second wave in the winter of 2020/21 is larger, as well as the assumption that infection rates rise as vaccinations proceed, as described at the beginning of this section.

vaccines are gradually deployed, but the path for mitigation remains at the estimated empirical benchmark path. Note that this path implies very modest shutdowns for 2021. In the second scenario, vaccines are distributed and mitigation is chosen optimally (the red line in Figure 6). In the third scenario, no vaccines are distributed, but the path for mitigation is chosen optimally (the blue line in Figure 6). The rows report five outcomes. These are the welfare gains for a utilitarian government, the expected welfare gain for older households, the number of deaths avoided relative to the benchmark scenario, the increase in total GDP for 2021, and the change in the debt-to-GDP ratio relative to the benchmark

[Table 5 about here]

The first column indicates that welfare gains from vaccines are equivalent to increasing consumption permanently by 0.72 percent for individuals' remaining lifetimes.⁵⁷ The gains for the old are much larger, at 6 percent of consumption. Vaccines, holding mitigation policy constant, save 193,000 lives. On the other hand, the impact on GDP from having fewer workers out sick is positive but small, with a boost to GDP in 2021 of 0.14 percent. With fewer people sick and receiving transfers, less government debt is accumulated.

The second column of the table shows the gains from the most favorable scenario: a vaccine is deployed, and the path for mitigation is chosen optimally to be more stringent in early 2021. Now, relative to the vaccine-only scenario, an extra 121,000 lives are saved, although at the cost of a nontrivial recession in the spring of 2021 and an associated deterioration of the fiscal position of the government.

The third column shows that *absent* a vaccine, the government can attain a substantial share of the welfare gains from the vaccine only scenario (column 1) if it moves from the baseline (very modest) mitigation path in 2021 to the optimal one plotted in Figure 6. However, the source of welfare gains is different: an extensive and persistent path for mitigation saves even more lives (264,000) but at the cost of a prolonged recession in 2021 in which output falls by more than 6 percent and the debt-to-GDP ratio rises by two and a half percentage points.

Thus, while a vaccine is unsurprisingly beneficial for welfare, it remains important to choose an appropriate path for economic mitigation and to adjust fiscal policy accordingly. That mitigation path (Figure 6) involves a hard lockdown to save lives in the short run, coupled with rapid

⁵⁷Given U.S. consumption of around \$15Tr, this corresponds to \$108bn per year on a flow basis. For comparison, the (one time) budget for Operation Warpspeed was \$10bn.

relaxation to save the economy from a prolonged recession as vaccinations progress and infections slow.

7 Conclusion

A key challenge in designing an optimal policy response to a pandemic is that mitigation efforts offer large potential benefits to some groups (in the case of COVID-19, they benefited the old) while imposing large costs on others, typically young participants in the labor market, especially those working in sectors directly affected by a lockdown. The fact that the gains and losses from mitigation are unequally distributed makes fighting pandemics politically 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, when redistribution creates additional distortions and thus is socially costly, mitigation policy may not be Pareto-improving in practice.

Within a quantitative heterogeneous agent model with government debt and micro-founded distortions from redistribution, we have computed the optimal joint shutdown and redistribution policy, in both the absence and the presence of a vaccination campaign. Relative to predictions from models without heterogeneity or with costless redistribution, the utilitarian optimal shutdown is milder, at the cost of more adverse health outcomes. Our baseline calibration suggests that, relative to what a utilitarian government would choose, the shutdown actually in place in the spring of 2020 was slightly too extensive, but was lifted too quickly. The presence of government debt allows the government to smooth the distortions generated by the need for redistribution over time, and results in more substantial shutdowns than in the absence of debt.

When a vaccine becomes available that actually saves lives, rather than simply postponing deaths, optimal mitigation in the second wave of infections in January becomes stronger, but much shorter, relative to that in a world in which the vaccine is not available. Extrapolating these lessons across countries, and for past and future pandemics, we conclude that different regions of the world should pursue rather different policies. The West and richer countries in Asia are first in line for vaccines, have access to sovereign debt markets and well-developed institutions for social insurance and redistribution, and have large population shares of vulnerable old people. Such countries should shut down hard early and then open up quickly when a majority of the population (and especially among the old) have been vaccinated. In contrast, much of the developing world is last in line for vaccines, has limited fiscal capacity to implement redistribution,

and has relatively few elderly residents. For such countries, our model suggests that much more limited shutdowns will be optimal.

By the same token, past pandemics might offer only limited insights for the future. In 1918 there was no realistic prospect for the speedy development of a vaccine against the Spanish flu, and no countries had extensive public insurance systems. In such a context, extensive lockdown policies would have come with high costs and uncertain benefits. On the other hand, if the current research into novel vaccines makes vaccine development against current or future virus mutations (such as “Omicron”) or other diseases even more rapid, our work suggests that a government with strong institutions and fiscal space should lock down hard and early and then open up quickly as the vaccination campaign progresses.

References

- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. WHINSTON (2021): “Optimal Targeted Lockdowns in a Multi-Group SIR Model,” American Economic Review: Insights, 3, 487–502.
- ARELLANO, C., Y. BAI, AND G. MIHALACHE (2023): “Deadly Debt Crises: COVID-19 in Emerging Markets,” Review of Economic Studies, forthcoming.
- ARGENTE, D., F. ALVAREZ, AND F. LIPPI (2021): “A Simple Planning Problem for COVID-19 Lockdown, Testing, and Tracing,” American Economic Review: Insights, 3, 367–382.
- ARIAS, E. AND J. XU (2020): “United States Life Tables, 2018,” National Vital Statistics Reports, 69, 1–44.
- ATKESON, A. (2020): “On Using SIR Models to Model Disease Scenarios for COVID-19,” Federal Reserve Bank of Minneapolis Quarterly Review, 41, 1–35.
- BAJEMA, K. L. ET AL. (2021): “Estimated SARS-CoV-2 Seroprevalence in the US as of September 2020,” JAMA Internal Medicine, 181, 450–460.
- BAYER, C. AND M. KUHN (2020): “Intergenerational Ties and Case Fatality Rates: A Cross-Country Analysis,” IZA Discussion Papers 13114, Institute of Labor Economics (IZA).
- BENABOU, R. (2002): “Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?” Quarterly Journal of Economics, 70, 481–517.
- BERGER, D., K. HERKENHOFF, C. HUANG, AND S. MONGEY (2022): “Testing and Re-opening in an SEIR model,” Review of Economic Dynamics, 43, 1–21.

- BOGNANNI, M., D. HANLEY, D. KOLLINER, AND K. MITMAN (2020): "Economics and Epidemics: Evidence from an Estimated Spatial Econ-SIR Model," London, Centre for Economic Policy Research, DP 15329.
- BOPPART, T., K. HARMENBERG, J. HASSLER, P. KRUSELL, AND J. OLSSON (2022): "Integrated Epi-Econ Assessment of Vaccination," Journal of Economic Dynamics and Control, 140.
- BRAZEAU, N. F. ET AL. (2020): "Report 34: COVID-19 Infection Fatality Ratio: Estimates from Seroprevalence," Imperial College COVID-19 response team.
- 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," CRC TR 224 Discussion Paper Series, University of Bonn and University of Mannheim.
- BUITRAGO-GARCIA, D. ET AL. (2020): "Occurrence and Transmission Potential of Asymptomatic and Presymptomatic SARS-CoV-2 Infections: A Living Systematic Review and Meta-analysis," PLOS Medicine, 17, 1–25.
- CDC (2020): "COVID-19 Pandemic Planning Scenarios," Centers for Disease Control and Prevention, <https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios-archive/planning-scenarios-2021-03-19.pdf>.
- CHARI, V., R. KIRPALANI, AND C. PHELAN (2021): "The Hammer and the Scalpel: On the Economics of Indiscriminate versus Targeted Isolation Policies during Pandemics," Review of Economic Dynamics, 42, 1–14.
- EICHENBAUM, M., S. REBELO, AND M. TRABANDT (2021): "The Macroeconomics of Epidemics," The Review of Financial Studies, 34, 5149–5187.
- ENGLE, S., J. KEPPO, M. KUDLYAK, E. QUERCIOLO, L. SMITH, AND A. WILSON (2021): "The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics," Working Paper.
- FARBOODI, M., G. JAROSCH, AND R. SHIMER (2021): "Internal and External Effects of Social Distancing in a Pandemic," Journal of Economic Theory, 196.
- FERNANDEZ-VILLAVARDE, J. AND C. I. JONES (2022): "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities," Journal of Economic Dynamics and Control, 140.
- GARRIGA, C., R. E. MANUELLI, AND S. SIDDHARTHA (2022): "Optimal Management of an Epidemic: Lockdown, Vaccine and Value of Life," Journal of Economic Dynamics and Control, 140.
- GLOVER, A., J. HEATHCOTE, AND D. KRUEGER (2022): "Optimal Age-Based Vaccination and Economic Mitigation Policies for the Second Phase of the Covid-19 Pandemic," Journal of Economic Dynamics and Control, 140.

- GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J. V. RÍOS-RULL (2020a): “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic,” NBER Working Paper, No. 27046.
- GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RÍOS-RULL (2020b): “Intergenerational Redistribution in the Great Recession,” Journal of Political Economy, 128, 3730–3778.
- GOLLIER, C. (2021): “The Welfare Cost of Vaccine Misallocation, Delays and Nationalism,” Journal of Benefit-Cost Analysis, 12, 199–226.
- GONZALEZ-EIRAS, M. AND D. NIEPELT (2020): “On the Optimal “Lockdown” during an Epidemic,” London, Centre for Economic Policy Research, DP 14612.
- 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. <http://dx.doi.org/10.2139/ssrn.3561244>.
- GUERRIERI, V., G. LORENZONI, L. STRAUB, AND I. WERNING (2022): “Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?” American Economic Review, 112, 1437–1474.
- HALL, R. E., C. I. JONES, AND P. J. KLENOW (2020): “Trading Off Consumption and COVID-19 Deaths,” Federal Reserve Bank of Minneapolis Quarterly Review, 42.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2017): “Optimal Tax Progressivity: An Analytical Framework,” The Quarterly Journal of Economics, 132, 1693–1754.
- HONG, J. H., J. PIJOAN-MAS, AND J.-V. RÍOS-RULL (2018): “Health and Inequality,” Work in Progress. First Version, 2016.
- INSTITUTE FOR HEALTH METRICS AND EVALUATION (2020): “COVID-19 Results Briefing: the United States of America,” Available at http://www.healthdata.org/sites/default/files/files/Projects/COVID/briefing_US_20201119.pdf.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2020): “The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the U.S.” National Bureau of Economic Research Working Paper 27794.
- KRUEGER, D., H. UHLIG, AND T. XIE (2022): “Macroeconomic Dynamics and Reallocation in an Epidemic: Evaluating the Swedish Solution,” Economic Policy, 37, 341–398.
- LAUER, S. ET AL. (2020): “The Incubation Period of Coronavirus Disease 2019 (COVID-19) From Publicly Reported Confirmed Cases: Estimation and Application,” Annals of Internal Medicine, 172, 577–582.
- LUCAS, R. E. AND N. L. STOKEY (1983): “Optimal Fiscal and Monetary Policy in an Economy without Capital,” Journal of Monetary Economics, 12, 55–93.
- MCANDREW, T. (2020): “COVID19-Survey 8-2020-04-08,” Working Paper, University of Massachusetts Amherst.

- 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>.
- NEW YORK TIMES (2020): "Virus Knocks Thousands of Health Workers Out of Action in Europe," <https://www.nytimes.com/2020/03/24/world/europe/coronavirus-europe-covid-19.html>.
- O'DRISCOLL, M. ET AL. (2021): "Age-specific Mortality and Immunity Patterns of SARS-CoV-2," *Nature*, 590, 140–145.
- ROSSEN, L. ET AL. (2020): "Excess Deaths Associated with COVID-19, by Age and Race and Ethnicity — United States, January 26–October 3, 2020," *Morbidity and Mortality Weekly Report*, Centers for Disease Control and Prevention.
- SEPKOWITZ, K. (2020): "It's Shameful how Many Health-care Workers are Dying from Covid-19," *CNN Opinion*, April 15.
- TODA, A. A. (2020): "Susceptible-Infected-Recovered (SIR) Dynamics of COVID-19 and Economic Impact," Working Paper, UCSD.
- 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. JHA (2020): "American Hospital Capacity And Projected Need for COVID-19 Patient Care," *Health Affairs Blog*, March 17, 2020. DOI: 10.1377/hblog20200317.457910.
- WOOLF, S. H., D. A. CHAPMAN, R. T. SABO, ET AL. (2020): "Excess Deaths From COVID-19 and Other Causes, March-July 2020," *JAMA*, 324, 1562–1564.

Tables

Table 1: Epidemiological Parameter Values

Behavior-Contagion			
α_w	infection at work	35% of infections	0.25
α_c	infection through consumption	19% of infections	0.12
β_e	infection in hospitals	5% of infections at peak	0.80
β_h	infection at home	Initial R_0 of 2.5	0.10
$x^a(0)$	initial asymptomatic infections	deaths through April 12, 2020	578.23
Disease Evolution			
σ^{yaf}	rate for young asymptomatic into fever	69% fever, 5.1 days	$\frac{0.69}{5.1}$
σ^{yar}	rate for young asymptomatic into recovered		$\frac{0.31}{5.1}$
σ^{oaf}	rate for old asymptomatic into fever	69% fever, 5.1 days	$\frac{0.69}{5.1}$
σ^{oar}	rate for old asymptomatic into recovered		$\frac{0.31}{5.1}$
σ^{yfe}	rate for young fever into emergency	3.41% hospitalization, 7 days	$\frac{0.0341}{7}$
σ^{yfr}	rate for young fever into recovered		$\frac{0.966}{7}$
σ^{ofe}	rate for old fever into emergency	31.8% hospitalization, 7 days	$\frac{0.318}{7}$
σ^{ofr}	rate for old fever into recovered		$\frac{0.682}{7}$
σ^{yed}	rate for young emergency into dead	6.2% conditional mortality, 6.2 days	$\frac{0.062}{6.2}$
σ^{yer}	rate for young emergency into recovered		$\frac{0.938}{6.2}$
σ^{oed}	rate for old emergency into dead	26.6% conditional mortality, 8.1 days	$\frac{0.266}{8.1}$
σ^{oer}	rate for old emergency into recovered		$\frac{0.734}{8.1}$
Time Variation in Mortality			
δ	rate hospital mortality declines	30% decline over 6 months	0.71
ζ_H	scaling for transmission in winter	deaths to May 31 2020	0.56
ζ_L	scaling for transmission in summer	deaths to Oct 31 2020	0.47
T^s	date summer (low transmission season) starts	deaths to Dec 31 2020	April 22

Table 2: Economic Parameters

Preferences			
μ^y	share of young	85%	0.85
ρ	discount rate	3.0% per year	$\frac{0.03}{365}$
T^y	residual life expectancy young	47.8 years	365×47.8
T^o	residual life expectancy old	14.0 years	365×14.0
φ	utility weight on hours	normalization	1.0
χ	Frisch elasticity for hours	1.0	1.0
\bar{u}	value of life	VSL = 10.8× consumption p.c.	11.61
\hat{u}^f	disutility of fever	lose 30% of baseline utility	-3.24
\hat{u}^e	disutility of emergency care	lose 100% of baseline utility	-10.8
η	elasticity lux. demand to hospitalizations	CPI relative prices	-156.5
Technology and Fiscal Policy			
μ^b	size of basic sector	55%	0.55
g	pre-COVID govt. spending	24.7% of GDP	0.247
τ^*	pre-COVID tax rate	utilitarian optimal	0.303
T^*	pre-COVID transfer	budget balance	0.223
Θ	hospital capacity	100,000 beds	0.000303
λ_o	impact of overuse on mortality	25% higher mortality at 200,000	825

Table 3: Model Outcomes Through July 4, 2021

	Deaths	Infections	GDP Loss	Transfers	Debt
Baseline Mitigation	660,000	138.5m	9.8%	2.8%	3.5%
Utilitarian Mitigation	350,000	75.1m	13.7%	4.0%	4.7%
No Mitigation	1,100,000	136.8m	0.6%	0.1%	0.3%

Cumulative shortfalls in GDP, cumulative extra transfer payments, and debt on July 4 are all expressed as shares of pre-COVID annual GDP.

Journal Pre-proof

Table 4: Welfare Gains (+) or Losses (-) From Mitigation

Group	Policy				
	None	Old	Luxury	Basic	Utilitarian
Young Basic	0.04%	-0.23%	0.22%	0.23%	0.11%
Young Luxury	0.13%	-0.50%	0.30%	0.29%	0.04%
Old	-6.07%	6.67%	-1.39%	-0.38%	4.25%
Utilitarian	-0.44%	0.21%	0.12%	0.20%	0.42%

Journal Pre-proof

Table 5: Welfare Gains From Vaccine Introduction

	Vaccine Baseline Mitigation	Vaccine Optimal Mitigation	No Vaccine Optimal Mitigation
Utilitarian Welfare	0.72%	0.84%	0.53%
Old Welfare	6.06%	7.81%	6.35%
Deaths Avoided	193,000	314,000	264,000
GDP Gain, 2021	0.14%	-2.28%	-6.10%
Debt to GDP	-0.47%	0.92%	2.53%

Journal Pre-proof

Figures

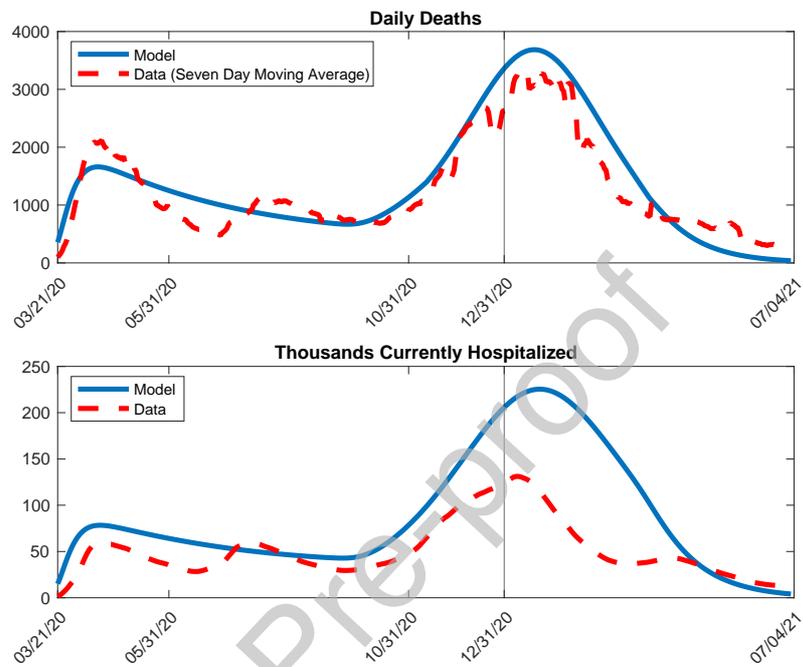


Figure 1: Daily Deaths and Current Hospitalizations

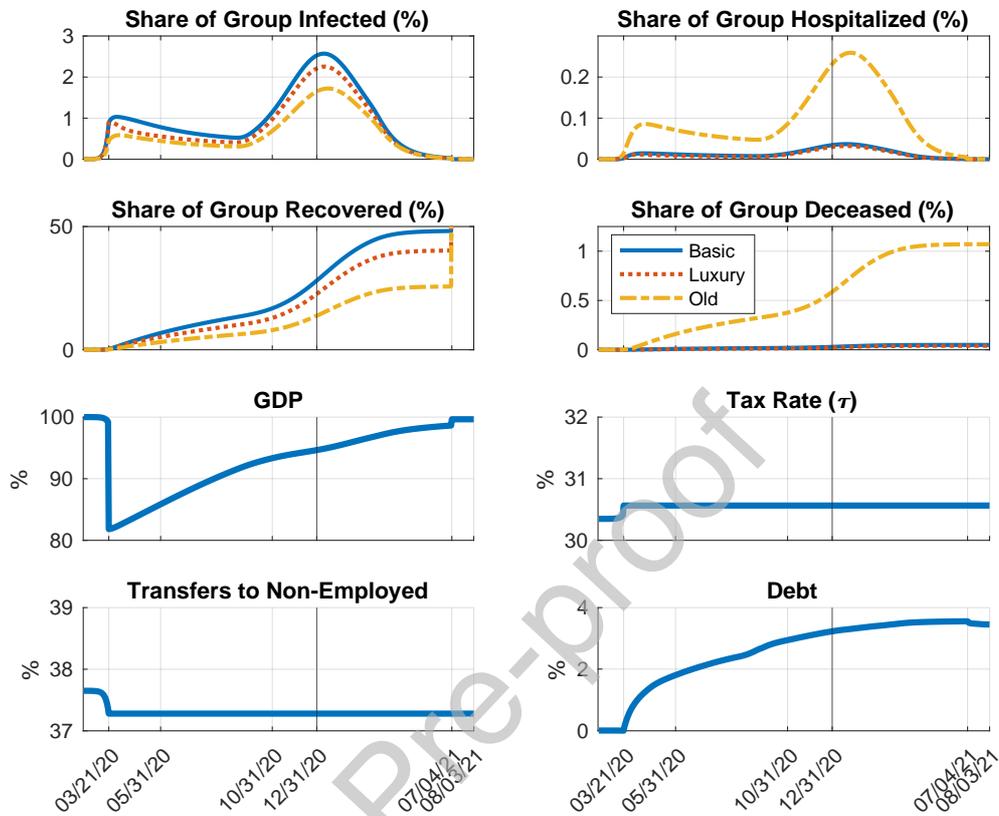


Figure 2: Upper 4 Panels: Health Outcomes: Shares of Each Group Infected, Hospitalized, Recovered, Deceased. Lower 4 Panels: Economic Outcomes: GDP, Tax Rate τ , Per Capita Transfers and Government Debt. GDP, Transfers and Debt are all Relative to Pre-COVID GDP.

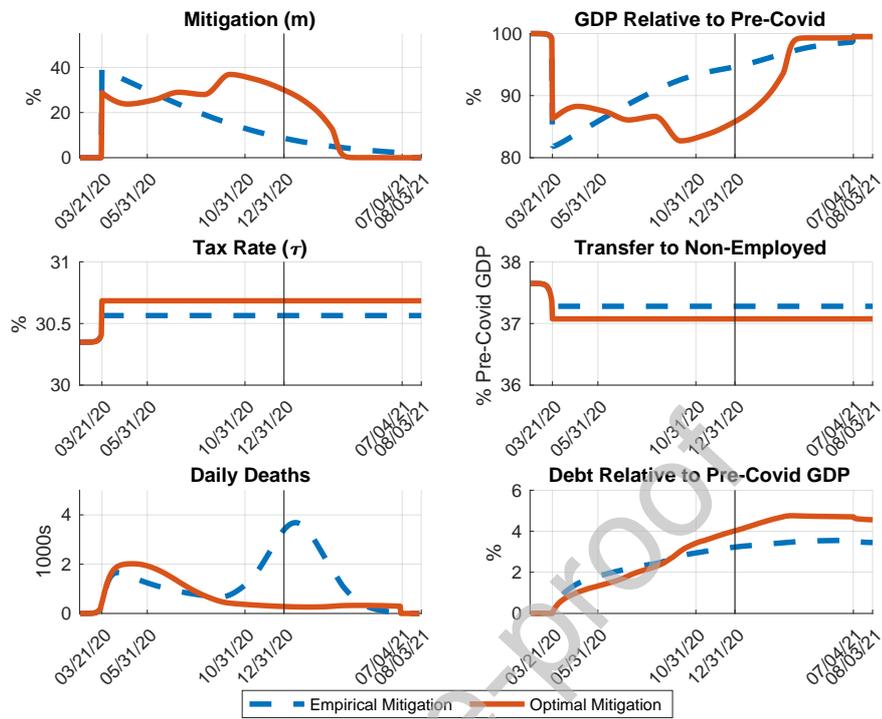


Figure 3: Economic and Health Outcomes: Utilitarian Optimal Mitigation versus Baseline Mitigation.

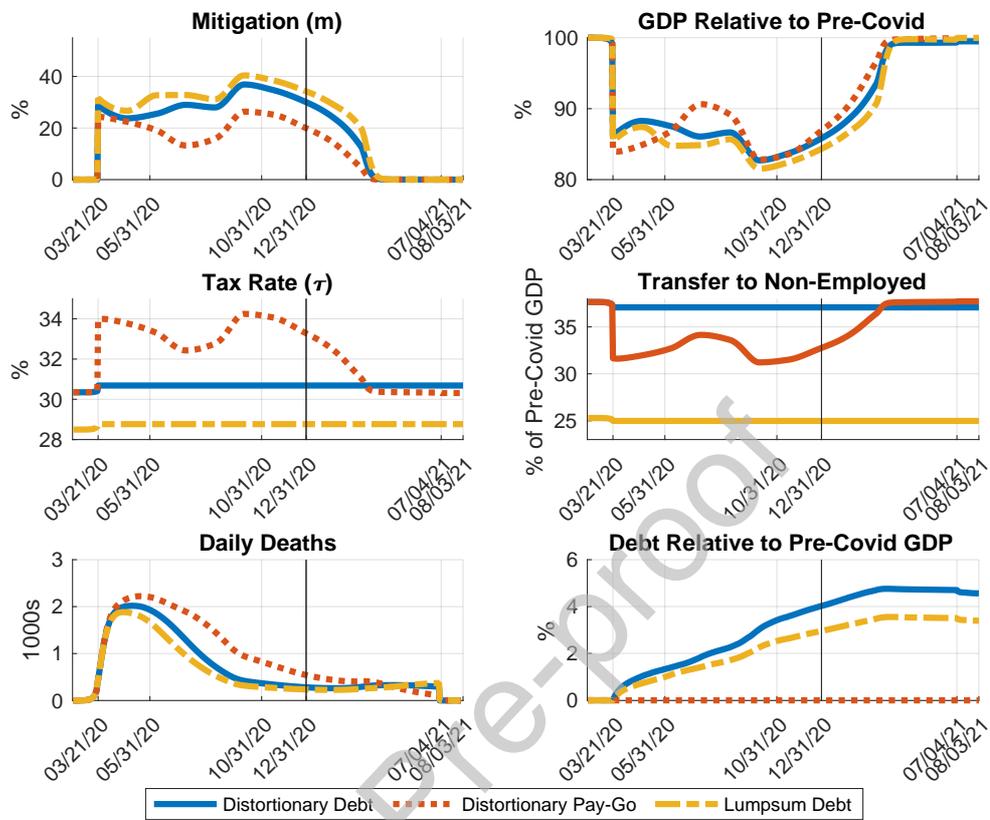


Figure 4: Preferred Mitigation Paths under Alternative Fiscal Rules: (1) Baseline Distortional Tax with Debt (Solid Blue), (2) Lump-Sum Taxes with Debt (Dashed Orange Lines), (3) Distortional Tax, No Debt (Dotted Red Lines)

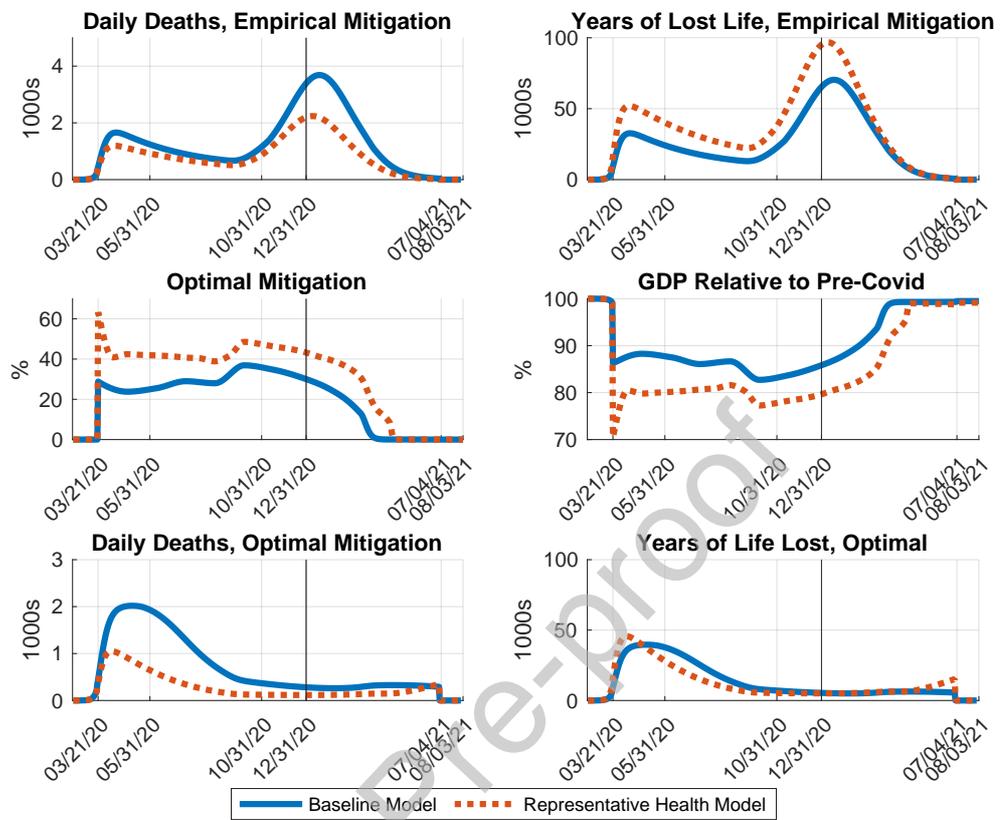


Figure 5: Role of Variation by Age in Health Threat from COVID-19: Baseline Model and Representative Health Model.

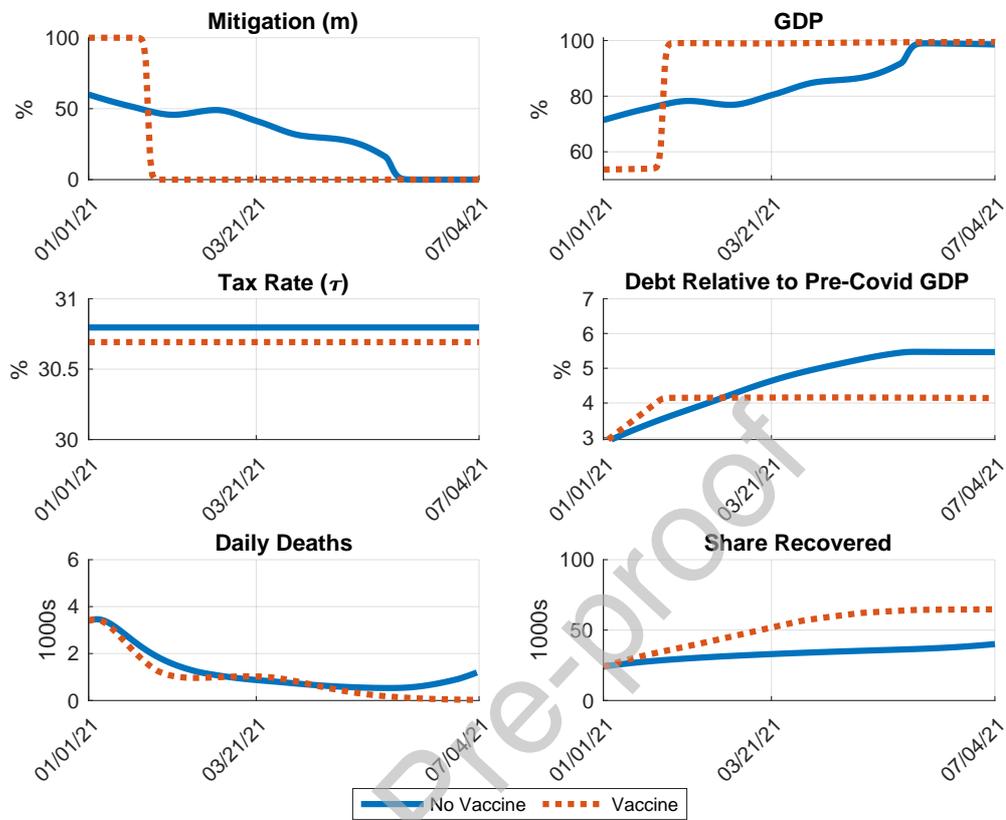


Figure 6: Utilitarian-Optimal Mitigation Paths with and without Vaccines