

DISCUSSION PAPER SERIES

DP14606

(v. 3)

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

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

LABOUR ECONOMICS

MACROECONOMICS AND GROWTH

MONETARY ECONOMICS AND FLUCTUATIONS

PUBLIC ECONOMICS

The logo for CEPR (Centre for Economic Policy Research) is a blue parallelogram with the letters 'CEPR' in white, bold, sans-serif font.

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

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

Discussion Paper DP14606
First Published 12 April 2020
This Revision 24 September 2021

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

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

- Labour Economics
- Macroeconomics and Growth
- Monetary Economics and Fluctuations
- Public Economics

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

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

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

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

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

Abstract

To to get the COVID-19 virus under control, many countries have shut down parts of the economy. Older individuals have the most to gain from slowing virus diffusion. Younger workers in sectors that are shuttered have most to lose. We build a model in which economic activity and disease progression are jointly determined. Individuals differ by age (young, retired), by sector (basic, luxury), and by health status. Disease transmission occurs in the workplace, through consumption, at home, and in hospitals. We study the optimal economic mitigation policy for a government that can redistribute through taxes and transfers, but where taxation distorts labor supply and output. Optimal redistribution and mitigation policies interact, and more modest shutdowns are optimal when redistribution creates tax distortions. A harder but shorter shutdown is preferred as vaccines become available in the first half of 2021.

JEL Classification: N/A

Keywords: COVID-19, Economic Policy, redistribution

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

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

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

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

Acknowledgements

We thank Josep Pijoan-Mas, Marc Suchard, Andrew Atkeson, Simon Mongey and Stefanie Stantcheva 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.

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

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

September 23, 2021

First version, March 31 2020

Abstract

To to get the COVID-19 virus under control, many countries have shut down parts of the economy. Older individuals have the most to gain from slowing virus diffusion. Younger workers in sectors that are shuttered have most to lose. We build a model in which economic activity and disease progression are jointly determined. Individuals differ by age (young, retired), by sector (basic, luxury), and by health status. Disease transmission occurs in the workplace, through consumption, at home, and in hospitals. We study the optimal economic mitigation policy for a government that can redistribute through taxes and transfers, but where taxation distorts labor supply and output. Optimal redistribution and mitigation policies interact, and more modest shutdowns are optimal when redistribution creates tax distortions. A harder but shorter shutdown is preferred as vaccines become available in the first half of 2021.

Keywords: COVID-19; Economic Policy; Redistribution

*We thank Josep Pijoan-Mas, Marc Suchard, Andrew Atkeson, Simon Mongey, Stefanie Stantcheva and four 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.

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

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

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

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

1 Introduction

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 centers on the question of how aggressively to restrict economic activity in order to slow the spread of a pandemic and how quickly to lift these 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 that the source of this disagreement is the fact that the benefits and costs of “economic 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 lockdown policies have vastly diverging preferences.

Standard epidemiological models describing the dynamic of a pandemic miss this disagreement because they assume a representative agent structure, in which all households face the same trade-off between restrictions on social interaction 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 bearing 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 will naturally focus on reducing activity in sectors in which there is a social aspect to consumption and sectors that produce goods or services perceived to be non-essential. For example, during the COVID-19 pandemic, 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.

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 2020, in which no vaccine was on the horizon, and then for the second phase of 2021, when a gradual roll out of an effective vaccine took place. Our model has the following 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 and transfers to redistribute the economic burden of the pandemic across individuals.

We distinguish between three types of people: young workers in a basic sector, young workers in a luxury sector, and old retired people. The output of workers in the two sectors is combined to produce a single final consumption good. Workers are immobile across sectors. 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 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: 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 co-consumers in the marketplace, from friends and family at home, and from the sick in hospitals. Because they do not work, the old do not face direct exposure at work, but virus transmission in the workplace indirectly increases infection rates in other settings. Our three different infected subgroups spread the virus in very different ways: the asymptomatic are unlikely to realize they are contagious and will continue to work and to consume; those with a fever will stay at home and 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 policy levers. First, at each date, it can choose what fraction of activity to shut down in the luxury sector. We call this policy 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 in the luxury sector. Second, the government chooses how much income to redistribute from those working to those who are not, because they are old, because they are unwell, or because their workplaces have been closed. Redistribution is desirable given a utilitarian social welfare function. But transfers must be financed by taxes that distort the labor-leisure choice of those working. Redistribution therefore imposes a deadweight loss on society, and perfect insurance between workers and non-workers is suboptimal. More mitigation is associated with higher tax rates, lower aggregate production and consumption, and more inequality across workers (which increases the more elastic is labor supply to marginal tax rates). The distortions induced by redistribution in turn reduce the

dynamic incentives of the government for mitigation. In particular, we prove that the marginal welfare costs of mitigation are larger when redistribution is costly than when the government has access to lump-sum taxes.

We use this model to characterize quantitatively the optimal path of mitigation, both for 2020 in the absence of a vaccine and, separately, for 2021, as vaccinations that protect individuals both from contracting and from spreading the disease are gradually administered. We first 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 wish to highlight three key 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 28 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, however, 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 longer and larger lockdowns than the young (72 percent versus 17 percent at the peak). 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 optimal one are about ten times larger for the old than the young, while welfare differences across the two young groups are relatively minor.¹

¹The 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

Second, the optimal mitigation path depends on the costs of redistribution through the tax system. If lump-sum taxes and transfers are available and redistribution is costless, it is easier to shut down the economy and compensate the losers. In contrast, if taxes distort economic activity, then weaker shutdowns are optimal. The optimal path with costly redistribution responds less to the waves of infections and deaths in 2020 and calls for less mitigation, especially in the peak of the second wave. The broader policy implication is that the optimal mitigation policy in response to a pandemic will be sensitive to the details of the social insurance system and to the amount of fiscal space that a country enjoys.

Third, 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), 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 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 more approximately 65 percent of the luxury sector in the second wave of January 2021, when it knows effective vaccines will be distributed, whereas the lockdown would be 10 percentage points smaller in that wave if there were no vaccine on the horizon. 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: by the end of March, the locked-down share of the luxury sector is 10 percent lower with a vaccine than without. The broader lesson we draw from this analysis is that vaccines and lockdowns are policy complements rather than substitutes, at least in the short run, giving some support to the idea that one has to beat the virus first in order to subsequently fix the economy.

the preferences of the young and to therefore be too modest relative to the utilitarian optimum.

1.1 Related Literature

Our paper contributes to a rapidly expanding 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. (2020) and Argente et al. (2021).²

We wish to emphasize 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. (2020) and Brotherhood et al. (2020).

A second distinctive feature of our framework is that the economic side of our model is modeled in an explicit structural way. Each of the key household constituencies solves an explicit maximization problem subject to a budget constraint, and there is no fictitious representative household that pools the economic costs of shutdowns. As a result 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 economic model of Kaplan et al. (2020) features even 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 we characterize the optimal redistribution policy via taxes and transfers in a setting in which we can solve for equilibrium economic allocations

²Other notable papers in the rapidly growing economics literature on COVID-19 include Fernandez-Villaverde and Jones (2020), Greenstone and Nigam (2020), Krueger et al. (2020), Toxvaerd (2020), Farboodi et al. (2021), Guerrieri et al. (2020), Bayer and Kuhn (2020), Berger et al. (2020), Chari et al. (2021), Hall et al. (2020) and Toda (2020).

in closed form, as a function of the current health distribution and the extent of mitigation. We can also characterize optimal taxes and transfers in closed form and show how they vary with the share of the population that relies on government support. This closed-form model of optimal redistribution could be applied to other contexts.

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., 2021). 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.

The paper unfolds as follows. In Section 2, we start by describing how we model the joint evolution of the economy and the population. In Section 3, we then turn to describe how we model mitigation and redistribution policies and how we go about solving for optimal policies. The calibration strategy is described in Section 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 first describe the individual state space, setting out the nature of heterogeneity by age and health status. In Section 2.2, we describe the multi-sector production technology and how mitigation shapes the pattern of production. Section 2.3 explains the details of our *SAFER* extension of the standard *SIR* epidemiological model and the channels of disease transmission.

2.1 Household Heterogeneity

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. Technology 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 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.³ 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 symptoms. Those with a fever are sufficiently sick to know they are likely contagious, and they stay at home and avoid the workplace and market consumption. Those requiring emergency care are hospitalized. The recovered are again healthy, no longer contagious, and immune from future infection. A worst-case virus progression is from susceptible (S) to asymptomatic (A) to fever (F) to emergency care (E) to dead (D).⁴ However, recovery (R) is possible from the asymptomatic, fever, and emergency-care 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 the basic b sector, while the rest of the young of mass x^ℓ work in a luxury sector, denoted ℓ . We assume that only luxury consumption is a potential source of COVID infection and that shutdowns apply only to the luxury sector. In particular, shutdowns will require some or all of the workers in the ℓ sector to stay at home in order to reduce virus transmission in the workplace and through luxury consumption. We call such a

³Thus, 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. (2020).

⁴Note 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. (2020) make a similar modeling choice.

policy a mitigation policy, m . More precisely, m_t denotes the fraction of luxury workers who are instructed to not go to work at time t . We assume that workers cannot change sectors (at least, not during the short time horizon studied in this paper); thus, the sector of work is a fixed characteristic of a young individual. In terms of notation, superscripts denote the dimensions of household heterogeneity, 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.

We assume that 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.⁵

In contrast to the basic sector, output in the luxury sector depends on mitigation policy and is given by

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

Let the basic good be the numeraire, and let p denote the price of the ℓ good. GDP is given by

$$Y = Y^b + pY^\ell. \quad (3)$$

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

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

$t = 0$, $x^{yb} = \mu^y \mu^b$, $x^{y\ell} = \mu^y(1 - \mu^b)$, 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 original 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 as infection-generating rates $\beta_j(m_t)$.

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

$$x_w(m_t) = x^{yba} + (1 - m_t)x^{y\ell a}, \quad (4)$$

$$x_c = x^a, \quad (5)$$

$$x_h = x^a + x^f, \quad (6)$$

$$x_e = x^e, \quad (7)$$

where these numbers reflect the assumption that symptomatic individuals neither go to work nor go shopping and that basic and luxury sector workers meet in the workplace. Note that the number of contagious workers depends on the mitigation parameter 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.⁶ 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, to be calibrated below, governing the intensity by which meetings among individuals 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 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, and the economy-wide infection-generating rate for work-related infections is therefore given by the following employment-share-weighted average:

$$\beta_w(m_t) = \frac{x^{bw}}{x^{bw} + (1 - m_t)x^{\ell w}} \times \alpha_w + \frac{(1 - m_t)x^{\ell w}}{x^{bw} + (1 - m_t)x^{\ell w}} \times \alpha_w(1 - m_t). \quad (8)$$

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,

$$\beta_c(m_t) = \frac{(1 - m_t)x^{\ell w}}{(1 - \mu^b)\mu^y} \times \alpha_c(1 - m_t), \quad (9)$$

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 $\beta_c(0) = \alpha_c$.

⁶See Xu et al. (2020) for more detailed evidence on infection patterns in the workplace.

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

$$\dot{x}^{ybs} = -[\beta_c(m_t)x_c + \beta_h x_h] x^{ybs} - \beta_w(m_t)x_w(m_t) x^{ybs} - \beta_e x_e x^{ybs}, \quad (10)$$

$$\dot{x}^{y\ell s} = -[\beta_c(m_t)x_c + \beta_h x_h] x^{y\ell s} - \beta_w(m_t)x_w(m_t)(1 - m_t)x^{y\ell s}, \quad (11)$$

$$\dot{x}^{os} = -[\beta_c(m_t)x_c + \beta_h x_h] x^{os}. \quad (12)$$

Consider the first outflow rate in equation (10). 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 (6). 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 (10), equation (11) 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, workers in the luxury sector do not take care of sick patients in hospitals, and thus the last term in equation (10) is absent in equation (11). Equation (12) 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 (33) to (44) 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

$$\varphi(x^e) = \lambda_o \max\{x^e - \Theta, 0\}. \quad (13)$$

In (13) 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

$$E_0 \left\{ \int_0^{T^y} e^{-\rho t} S_t^y [u(c_t^y, h_t^y) + \bar{u} + \widehat{u}_t^i] dt \right\}, \quad (14)$$

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 , \widehat{u}_t^i . We will assume that $\widehat{u}_t^s = \widehat{u}_t^a = \widehat{u}_t^r = 0$ and that $\widehat{u}_t^e < \widehat{u}_t^f < 0$. 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 .

Note that workers who have a fever or are in hospital do not work. Neither does a fraction m of luxury sector workers whose workplaces have been shut down by mitigation policy. In equilibrium, expected utility of a young individual will depend on the sector in which she works, for three 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. Third, the relative demand for basic goods versus luxuries, and thus relative offered wages in the two sectors, will vary owing to pandemic-induced fear of infection, as we will shortly describe.

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 $c = c^b + c^\ell(1 - \xi_c)$ is a linear aggregate of consumption in the two sectors, and ξ_c is a negative preference shock to the consumption of luxury goods. This preference shock will be a function of the state of the pandemic in the full dynamic model, but for the exposition and derivation of static social welfare, we can treat it as a parameter. It will capture the shift in demand away from luxury goods, such as restaurant meals, as the pandemic unfolds.

Since one unit of labor produces one unit of output, and the basic good is the numeraire, the wage of sector b workers is $w^b = 1$. Similarly, one unit of time in the luxury sector produces output worth p , and perfect competition in the labor market implies that the luxury wage is $w^\ell = p$. 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. In any equilibrium in which both goods are produced and consumed in positive amounts, the price of luxury goods must satisfy

$$p = 1 - \xi_c, \tag{15}$$

which is the only price at which households are indifferent between consuming both goods. Healthy, non-mitigated households solve, for $i \in \{b, \ell\}$ with wages $w^b = 1$, $w^\ell = 1 - \xi_c$, the

maximization problem

$$\begin{aligned}\max_{c,h,c^b,c^\ell} U^i &= \log\left(c - \frac{h^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}}\right), \\ c &= c^b + c^\ell(1 - \xi_c), \\ c^b + (1 - \xi_c)c^\ell &= (1 - \tau)h w^i,\end{aligned}$$

with solution

$$\begin{aligned}h^i &= [(1 - \tau)w^i]^\chi, \\ c^i &= (1 - \tau)h^i w^i = [(1 - \tau)w^i]^{1+\chi}, \\ U^i &= -\log(1 + \chi) + (1 + \chi)\log(1 - \tau) + (1 + \chi)\log(w^i).\end{aligned}$$

Period utility of b workers is unaffected by the preference shock ξ_c to the ℓ goods. Good ℓ consumption is valued less when ξ_c goes up, but the price of ℓ goods also falls, and the two effects cancel out. In contrast, ℓ workers also see their wage $w^\ell = 1 - \xi_c$ fall, resulting in a utility loss for them. For non-working households, the budget constraint and period utility are

$$\begin{aligned}c^b + c^\ell(1 - \xi_c) &= c = T, \\ U^n &= \log(c) = \log(T).\end{aligned}$$

2.6 Aggregation and Market Clearing

The government raises taxes at rate τ , pays transfers T (in terms of the basic good), and purchases goods of both sectors of the economy. The market clearing conditions are

$$\begin{aligned}C^b + G^b &= x^{bw} h^b \equiv Y^b, \\ C^\ell + G^\ell &= (1 - m)x^{\ell w} h^\ell \equiv Y^\ell.\end{aligned}$$

We assume that government purchases are a constant share g of output in both sectors: $G^i = gY^i$ for $i \in \{b, \ell\}$. Aggregating across all households and exploiting the market clearing

conditions yields

$$\begin{aligned} C^b + G^b &= Y^b = x^{bw} h^b = x^{bw} [1 - \tau]^\chi, \\ C^\ell + G^\ell &= Y^\ell = (1 - m)x^{\ell w} h^\ell = (1 - m)x^{\ell w} [(1 - \tau)(1 - \xi_c)]^\chi. \end{aligned}$$

For future reference, we write the measures of working and nonworking households as

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

3 Fiscal Policy: Taxes and Transfers

In addition to financing public consumption parameterized by g , the government is responsible for two choices: a path of mitigation (shutdowns) m_t , and redistribution through proportional taxation on workers at rate τ , which finances lump-sum transfers T to individuals who do not or cannot work. These include those unemployed because of shutdowns, those with fever or who are hospitalized, and those who have retired. We assume that when setting taxes and transfers, the government values all individuals the same, implying equal Pareto weights in the social welfare function.⁸

Notice that there are no dynamic consequences of the tax-transfer policies. In particular, this part of government policy has no impact on any health transitions. At each date t we can therefore solve a static optimal tax-transfer policy problem that maximizes instantaneous social welfare given the current level of mitigation $m_t = m$ and the current population health distribution. We denote this maximum by $W^*(m)$.

⁸That the labor income of all workers is taxed at a constant and identical rate (rather than a potentially sector-specific and general progressive schedule) is an exogenous restriction of the policy space. In contrast, given the utilitarian social welfare function and a period utility function that is separable in consumption and health (including the value of being alive) it is in fact optimal to equalize transfers T across all nonworking individuals.

Government spending, taxes and transfers are restricted by the budget constraint

$$x^n(m)T + gY^b + pgY^\ell = \tau [x^{bw}w^bh^b + (1-m)x^{\ell w}w^\ell h^\ell] = \tau [Y^b + pY^\ell(m)] = \tau Y(m, \tau), \quad (16)$$

and thus per capita transfers to non-working households are given by

$$T = \frac{(\tau - g) [Y^b + pY^\ell(m)]}{x^n(m)} = \frac{(\tau - g)}{x^n(m)} (1 - \tau)^\chi Y^{LS}(m), \quad (17)$$

where $Y^{LS}(m)$ defines what aggregate output would be if taxes were lump-sum (i.e., with $\tau = 0$):

$$Y^{LS}(m) = x^{bw} + (1-m)x^{\ell w}(1 - \xi_c)^{1+\chi}. \quad (18)$$

Also note that the elasticity of tax revenues $\tau Y(m, \tau)$ with respect to the tax rate τ , given optimal labor supply of households, is given by

$$\epsilon_\tau \equiv \frac{\partial \tau Y(m, \tau)}{\partial \ln \tau} = 1 - \frac{\chi \tau}{1 - \tau}, \quad (19)$$

and thus the marginal excess burden of taxes, which measures the resources lost through behavioral responses per dollar of revenue raised, is given by $MEB(\tau) = \frac{1}{\epsilon_\tau} - 1 = \frac{\chi \tau}{1 - \tau - \chi \tau}$. The marginal excess burden of taxation is strictly increasing both in the level of the tax rate τ as well as in the size of the labor supply distortion parameterized by χ . We will soon prove that increasing mitigation implies higher optimal tax rates to finance transfers to a larger pool of non-workers. Thus increasing mitigation increases the excess burden of taxation, the more so the more elastic is labor supply.⁹

To determine static social welfare, note that period utilities for non-working individuals and for those working in the basic and luxury sectors are given by

$$U^n(\tau, m) = \log(T) = \log(\tau - g) + \chi \log(1 - \tau) + \log \left[\frac{Y^{LS}(m)}{x^n(m)} \right], \quad (20)$$

$$U^b(\tau) = -\log(1 + \chi) + (1 + \chi) \log(1 - \tau), \quad (21)$$

⁹In our calibration, $\chi = 1$ and pre-COVID $\tau^* = 0.303$, implying a marginal excess burden of 0.77. As τ^* rises to 0.36 when shutdowns begin, this burden rises to 1.29.

$$U^\ell(\tau) = -\log(1 + \chi) + (1 + \chi) \log(1 - \tau) + (1 + \chi) \log(1 - \xi_c) = U^b + (1 + \chi) \log(1 - \xi_c). \quad (22)$$

3.1 The Static Optimal Policy Problem

For a given current level of mitigation $m_t = m$, and taking as given the population measures, the objective of the utilitarian government is to maximize, by choice of the tax rate τ and the implied transfer T , average static utility from the corresponding consumption-labor allocation, which amounts to the following social welfare expression:

$$W(\tau, m) = x^{bw} U^b(\tau) + (1 - m)x^{\ell w}(m) U^\ell(\tau) + x^n(m) U^n(m, \tau) = \Gamma(\tau, m) + \Psi(m), \quad (23)$$

where

$$\Gamma(\tau, m) = (x^w(m) + \chi x^n(m)) \log(1 - \tau) + x^n(m) \log(\tau - g), \quad (24)$$

$$\Psi(m) = x^w(m) \log\left(\frac{1}{1 + \chi}\right) + (1 - m)x^{\ell w}(1 + \chi) \log(1 - \xi_c) + x^n(m) \log\left[\frac{Y^{LS}(m)}{x^n(m)}\right]. \quad (25)$$

Note that $\Psi(m)$ is a function solely of m and thus irrelevant for maximization with respect to τ . The optimal proportional tax rate can be found by setting the first order condition with respect to $\Gamma(\tau, m)$ to zero, which yields the following proposition, proved in Appendix B.1.

Proposition 1 *For given m , the optimal redistribution policy is given by*

$$\tau^*(m) = (1 - g) \frac{\mu^n(m)}{1 + \chi} + g, \quad (26)$$

$$T^*(m) = \left(\frac{1 - g}{1 + \chi}\right)^{1 + \chi} [\chi + \mu^w(m)]^\chi \frac{Y^{LS}(m)}{x}, \quad (27)$$

where $\mu^n(m) = x^n(m)/x$ is the share of the population that is not working and $\mu^w(m) = x^w(m)/x$ is the share of the population that is working. The optimal tax rate is strictly increasing in mitigation m and the optimal per capita transfer is strictly decreasing in m . The more elastic is labor supply (the larger is χ), the smaller are (τ^*, T^*) . The optimal tax rate is strictly increasing in the government expenditure share g , and the optimal transfer is strictly decreasing in g .

The key takeaway from this proposition is that more mitigation translates into higher optimal tax rates. These higher tax rates translate into larger distortions, which, as we show next, amplify the negative effect of mitigation m on static welfare. This additional cost of mitigation would be absent in a hypothetical alternative economy in which resources could be redistributed to non-working households in a non-distortionary fashion.

3.2 Evaluating Static Social Welfare

Substituting the expressions for τ^* and T^* in Proposition 1 into the period utilities (equation 20), one can show that static utility from consumption net of the disutility from labor of all three groups depends on mitigation only through the share of the working population, $\mu^w(m)$:

$$\begin{aligned}
 U^n[\mu^w(m)] &= (1 + \chi) \log \left(\frac{1 - g}{1 + \chi} \right) + \chi \log (\chi + \mu^w(m)) + \\
 &\quad \log \left[(1 - \xi_c)^{1+\chi} \left(\frac{\chi^{bw}}{\chi} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m) \right) \right], \\
 U^b[\mu^w(m)] &= (1 + \chi) \log \left(\frac{1 - g}{1 + \chi} \right) + (1 + \chi) \log (\chi + \mu^w(m)) - \log (1 + \chi), \\
 U^\ell[\mu^w(m)] &= U^b(\mu^w(m)) + (1 + \chi) \log(1 - \xi_c). \tag{28}
 \end{aligned}$$

For those working, more mitigation—and thus a lower share of workers $\mu^w(m)$ —is costly because it means a higher optimal tax rate, while for non-workers, it is costly because it means lower per capita transfers. Thus, for all types, static utility $U^i(\mu^w(m))$ is reduced. To establish that aggregate static social welfare defined in equation (23) is strictly decreasing in mitigation m , it is then sufficient to show that utility of non-mitigated ℓ -workers is higher than that of mitigated ℓ -workers (i.e., the non-workers), $U^\ell(\mu^w(m)) > U^n(\mu^w(m))$ for all m , so that a shift in the mass of people from ℓ workers to non-workers unambiguously lowers overall social welfare. The following assumption, stated purely in terms of the fundamentals of the economy, guarantees that this is the case even for $m = 0$ and thus for all $m \in [0, 1]$.

Assumption 1 *The preference shock ξ_c satisfies*

$$\xi_c < \bar{\xi}_c = 1 - \left[\frac{\chi(1 - \mu^w(m=0))}{(1 + \chi)\mu^{bw}} + 1 \right]^{\frac{-1}{1+\chi}} \in (0, 1). \quad (29)$$

The negative preference shock cannot be too large, because a larger shock would depress luxury sector prices and wages below the level at which ℓ -workers prefer to work rather than be mitigated. Define static social welfare from the consumption allocation as $W^*(m) = W(\tau^*, m)$; see equation (23).

Proposition 2 *For a given health distribution \mathbf{x} , define static social welfare in economies with distortionary taxes and with lump-sum taxes as $W^*(m)$ and $W^{LS}(m)$, respectively. Let Assumption 1 be satisfied. Then,*

1. $W^*(m)$ and $W^{LS}(m)$ can be decomposed as

$$\begin{aligned} W^*(m) &= x \widetilde{W}^*(\mu^w(m)), \\ W^{LS}(m) &= x \widetilde{W}^{LS}(\mu^w(m)), \end{aligned}$$

where per capita welfare \widetilde{W} depends on mitigation only through the share of workers $\mu^w(m)$ in the economy.

2. Per capita welfare $\widetilde{W}^*(\mu^w(m))$ and $\widetilde{W}^{LS}(\mu^w(m))$ is strictly increasing in $\mu^w(m)$ and thus strictly decreasing in mitigation m .
3. Per capita welfare satisfies

$$\frac{\partial \widetilde{W}^*(\mu^w(m))}{\partial \mu^w(m)} > \frac{\partial \widetilde{W}^{LS}(\mu^w(m))}{\partial \mu^w(m)} > 0.$$

This result states that static welfare is more sensitive to the share of workers and thus to mitigation m when taxes are distortionary relative to when taxes are lump-sum. Thus, the economic costs of mitigation are larger in the distortionary tax economy. With this characterization of static social welfare in hand, we now formulate the dynamic problem of the government.

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

$$\begin{aligned} \max_{m_t, \mathbf{x}_t} \int_0^{\bar{T}} e^{-\rho t} \left[W^*(\mathbf{x}_t, m_t) + x\bar{u} + \sum_{j \in \{yb, y\ell, o\}, i \in \{s, a, f, e, r\}} x^{ji} \widehat{u}^i \right] dt \\ + \frac{e^{-\rho \bar{T}}}{\rho} \left(1 - e^{-\rho(T^y - \bar{T})} \right) \left(x_{\bar{T}}^{yb} \left(U^b \left(\mu_{\bar{T}}^w(0) \right) + \bar{u} \right) + x_{\bar{T}}^{y\ell} \left(U^\ell \left(\mu_{\bar{T}}^w(0) \right) + \bar{u} \right) \right) \end{aligned} \quad (30)$$

$$+ \frac{e^{-\rho \bar{T}}}{\rho} \left(1 - e^{-\rho(T^o - \bar{T})} \right) x_{\bar{T}}^o \left(U^n \left(\mu_{\bar{T}}^w(0) \right) + \bar{u} \right), \quad (31)$$

where the share of the population working at the end of the COVID-19 planning horizon \bar{T} , given $m_{\bar{T}} = 0$, is $\mu_{\bar{T}}^w(0) = x_{\bar{T}}^y/x_{\bar{T}}$. The maximization is subject to the laws of motion of the health population distribution \mathbf{x}_t . The objective of the government has two components: welfare before time \bar{T} (first line) and continuation welfare after \bar{T} (second and third lines). The latter consists of the population-share-weighted residual lifetime utilities of the three groups, where after time \bar{T} there are no further utility-diminishing health state shocks. The former component reflects static social welfare from consumption-labor allocations $W^*(\mathbf{x}_t, m_t)$ (where we now make explicit the dependence on the health distribution \mathbf{x}_t), direct utility \bar{u} from being alive for the entire living population, and the negative utility impacts \widehat{u}^i of being sick with COVID-19.

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 $\frac{\partial W^*(\mathbf{x}, m)}{\partial m_t}$, characterized in Proposition 2, stemming from an increase in the tax rate and falls in output, consumption and per-capita transfers. 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 Parameter

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

Preferences We set the pure time discount rate in annual terms to 3 percent. We set the labor supply Frisch elasticity χ to one.

For the flow value of being alive, \bar{u} , we follow the value of a statistical life (VSL) approach. Greenstone and Nigam (2020) set the VSL to \$11.5 million, a figure based on work by the Environmental Protection Agency. The average age of Americans is 37.9, and at that age, remaining life expectancy is 42.6 years.¹⁰ If we discount at 3 percent, a one-time payment of \$11.5 million is equivalent to a flow payment for 42.6 years of \$478,000, which is 10.8 times yearly per capita consumption in the United States.¹¹

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

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

where \bar{c} is average per capita model consumption and r is the risk of death. Setting $VSL = 10.8\bar{c}$ and $r = 0$ gives $\bar{u} = 10.8 - \ln \bar{c}$. This implies an easily interpretable trade-off between mortality risk and consumption. For example, suppose we were to contemplate a shut-down that would reduce consumption for six months by 25 percent. By how much would this shut-down have to reduce mortality risk for an agent with 10 expected years of life for the agent to prefer the shutdown to no shutdown? The answer is the solution x to

$$\frac{1}{20} \ln(1 - 0.25) + \frac{19}{20} \ln(1) + 10.8 = (1 - x) (\ln 1 + 10.8),$$

which is 0.13%.

¹⁰We take life expectancy values from Table A of Arias and Xu (2020).

¹¹Per capita consumption in 2019 was \$44,272 : see <https://fred.stlouisfed.org/series/A794RCOA052NBEA>.

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 is equal to the flow value of being dead (zero). We think of typical young and old individuals as being 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, (2) shutdowns are assumed to be concentrated in the luxury sector, and (3) the demand for luxury consumption is inversely related to infection risk. To calibrate the relative employment and output shares of the two sectors, we partition the different components of consumer expenditure in the CPI into those that have a social aspect of consumption and those that do not. Thus, we 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 is a form of consumption involving social contact, because shutdowns have largely exempted healthcare. 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.¹²

Given $g = 0.247$ and $\chi = 1.0$, the optimal pre-COVID tax rate given by equation 26 is $\tau^* = 0.303$, implying a consumption level for workers of $c^b = c^\ell = 0.485$ and for the old of $c^o = T^* = 0.223$.

¹²Total 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).

Preference shifters We think of the preference shifter attached to luxury consumption, ξ_c , as depressing luxury demand when infection risk is high. Our interpretation is that social consumption is less enjoyable when you worry about potentially getting sick.¹³ We assume the following functional form:

$$\xi_{c,t} = 1 - \exp(\eta x_t^e),$$

where x_t^e is the number of people hospitalized – an imperfect but easily observable proxy for infection risk.

In the model equilibrium, p_t is exactly equal to $1 - \xi_{c,t}$. Thus, we estimate the parameter η by estimating the following regression:

$$\log p_t = \eta x_t^e + \varepsilon_t,$$

where the empirical counterpart for p_t is the expenditure-share weighted average monthly consumer price index for luxuries relative to basic goods, and x_t^e is hospitalizations as a share of population. The resulting estimate for η is -156.5 with a Newey-West standard error of 49.3 .¹⁴

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 will be 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,

¹³Note that this is an internal psychic cost. We have not modeled dynamic consumption choices in which individual consumption choices affect health and mortality risk in addition to delivering direct utility. Only the government, through its choice for m_t , can affect health outcomes in our model.

¹⁴Figure 12 in Appendix D visualizes the correlation between hospitalizations and the relative price of luxury goods in 2020.

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 assume this is common across age groups.¹⁵ 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.¹⁶

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, all else equal. Following Buitrago-Garcia et al. (2020) we assume that 31 percent of those infected recover without ever developing symptoms.¹⁷ 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 relative to people below that age. Second, even that differential understates the true impact of age on COVID lethality, since SARS-CoV-2 seroprevalence studies indicate that those over 65 have been infected at only around half the rate of younger individuals (see, for example, (Bajema

¹⁵Note 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.

¹⁶See Table 2 of CDC (2020).

¹⁷The 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.

et al., 2021)), something that will also be a feature of our model. A factor of 40 differential is broadly 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 $0.01/(0.85 + 0.15 * 40) = 0.146$ percent and an IFR for the old of $40 * 0.146 = 5.84$ percent.

Within our model the overall age-specific IFR is the product of the probability of becoming symptomatic, 0.69, times the age-specific probability that fever becomes sufficiently severe to require hospitalization, times the age-specific probability of death conditional on hospitalization. Of these 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 65.¹⁸ Given these deaths 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 $\beta_e = 0.80$.¹⁹ The values of α_w , α_c , and β_h determine the overall basic reproduction number R_0 value for COVID-19 and the share of disease transmission that occurs at work, via market consumption, and in non-market settings.

Mossong et al. (2008) find that 35% of potentially infectious inter-person contact happens in workplaces and schools, 19% occurs in travel and leisure activities, and the remainder takes place at home and in other settings.²⁰ 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

¹⁸The 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.

¹⁹5% 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).

²⁰Xu et al. (2020) discuss in detail heterogeneity in contact rates across different types of business (closed office, open office, manufacturing and retail) and a range of interventions that can reduce those rates.

and school transmission with transmission at work, travel and leisure with consumption-related transmission, and the residual categories with transmission at home. These targets are used to pin down choices for α_w and α_c , both relative to β_h , as follows.

The basic reproduction number R_0 is the number of people infected by a single asymptomatic person. For a single young person, assuming everyone else in the economy is susceptible and there is zero mitigation ($m = 0$), R_0^y is given by

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

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

The reproduction number expression for an old asymptomatic person, R_0^o , is similar, except that the old do not pass the virus on at work, but they are more likely to require hospitalization and transmit the virus in hospital:

$$R_0^o = \frac{\alpha_c + \beta_h}{\sigma^{oar} + \sigma^{oaf}} + \frac{\sigma^{oaf}}{\sigma^{oaf} + \sigma^{oar}} \frac{\beta_h}{\sigma^{ofr} + \sigma^{ofe}} + \frac{\sigma^{oaf}}{\sigma^{oaf} + \sigma^{oar}} \frac{\sigma^{ofe}}{\sigma^{ofe} + \sigma^{ofr}} \frac{\beta_e \mu^y \mu^b}{\sigma^{oer} + \sigma^{oed}}.$$

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

$$R_0 = \mu^y R_0^y + (1 - \mu^y) R_0^o. \quad (32)$$

The expected shares of workplace and consumption transmission are given by

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

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

Given these equations, we set the relative values α_w/β_h , α_c/β_h to replicate shares of workplace and consumption transmission equal to 35% and 19%. Note that this evidence does not pin down the *levels* of α_w , α_c , and β_h , or equivalently the level of R_0 . We set the level for these infection-generating-rate parameters to deliver (via equation 32) 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).

Hospital Capacity Tsai et al. (2020) estimate that 58,000 ICU beds are potentially available nationwide to treat COVID-19 patients. However, only 21.5% of COVID-19 hospital admissions require intensive care, suggesting that total hospital capacity is around $58,000/0.215=270,000$. Tsai et al. (2020) emphasize that this capacity is very unevenly allocated 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.²¹

4.2 Time Varying Parameters

Baseline Mitigation Our baseline path for mitigation, designed to approximate historical US 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

²¹Much 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.

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

model replicates the dynamics of employment from March through December 2020.²²

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 Remdesevir. 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 $\exp(-0.71 * 0.78) = 0.57$

²²We target the series for the Employment-Population Ratio [EMRATIO] published by the US 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 14 in the Appendix plots the implied model employment rate against the BLS series.

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. 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 will 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 will 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, which were 27,003, 108,102, 238,347 and 371,503, respectively.²³ We focus on matching deaths, since deaths are the best-measured and most welfare-relevant measure of the virus' impact.²⁴

²³All our data numbers for deaths are from the Centers for Disease Control: <https://covid.cdc.gov/covid-data-tracker/#datatracker-home>.

²⁴Following most of the existing literature, we focus on replicating official deaths, while recognizing that to the extent that the 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

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.²⁵ 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 on February 1.²⁶

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 0.56.²⁷ 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 will allow 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.

We assume that $\bar{T} = 1,000$ 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.²⁸

All epidemiological and economic parameter values are summarized in Tables 1 and 2.

This calibration implies the Spring 2020 population distribution by health status described in Table 3. On March 21, 0.82% of the US population was actively infected (including asymp-

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). One important open question is how many of these excess deaths reflect COVID illness versus how many might be attributed to deaths of despair tied to the recession and/or lockdowns (Mulligan, 2020).

²⁵There 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.

²⁶The precise timing of the first cases in the US remains unclear.

²⁷Recall 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 13 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).

²⁸This is sufficient time for the pandemic to have run its course in all our simulations.

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	47.8
T^o	residual life expectancy old	14.0 years	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

tomatic infections), with that number rising to 0.84% by April 12, with an additional 1.82% having recovered.²⁹

Table 3: Millions of People in Each Health State

	S	A	F	E	R	D (1,000s)
03/21/20	326.37	1.99	0.71	0.02	0.91	1.32
04/12/20	320.31	1.35	1.33	0.08	6.91	27.00

5 Looking Back: Optimal Mitigation During 2020

In this section, we describe the quantitative findings from our model. We start in Section 5.1 by comparing against data the model's predictions for health outcomes under the benchmark mitigation policy. We then explore its economic predictions in Section 5.2. We characterize

²⁹These numbers are within the range of expert estimates from the COVID-19 survey compiled by McAndrew (2020) at the University of Massachusetts.

the optimal mitigation path chosen by a utilitarian government in Section 5.3 and then show that the utilitarian optimal policy masks fundamental differences in policy preference across heterogeneous households (Section 5.4). Section 5.5 illustrates the importance of fiscal policy and the cost of redistribution for these results. After having validated the model against data from 2020, we then look forward and, in Section 6, characterize the optimal mitigation policy from January 1 2021 onward, taking into account the deployment of vaccines.

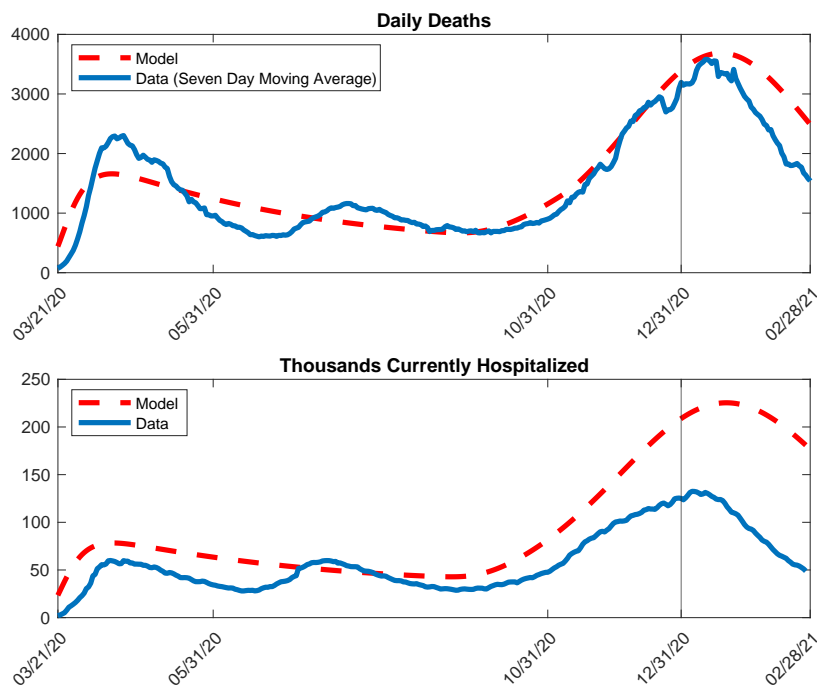


Figure 1: Model vs. Data: Daily Deaths and Current Hospitalizations.

5.1 The Health Pandemic through the Lens of the Benchmark Model

In Figure 1 we display U.S. time series for daily deaths and currently hospitalized individuals alongside their counterparts from the model simulation.³⁰ Observe that the model tracks actual

³⁰Both data series are from the Centers for Disease Control. The series for hospitalizations is at <https://covid.cdc.gov/covid-data-tracker/#hospitalizations> and the data for daily deaths is at https://covid.cdc.gov/covid-data-tracker/#trends_dailytrendscases.

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 somewhat 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, not all very sick infected individuals were hospitalized: some remained in nursing homes and other settings.

In Figure 2 we display the population health dynamics from March 21 to the end of February 2021. The top left panel shows the currently infected percentage of the population in each of the three groups (basic workers, luxury workers, and the old), and the top right panel plots the currently hospitalized share of each group. The bottom left panel depicts the total share of the population groups that has had COVID-19 and recovered, whereas the bottom right panel shows the population shares that have succumbed to the disease.

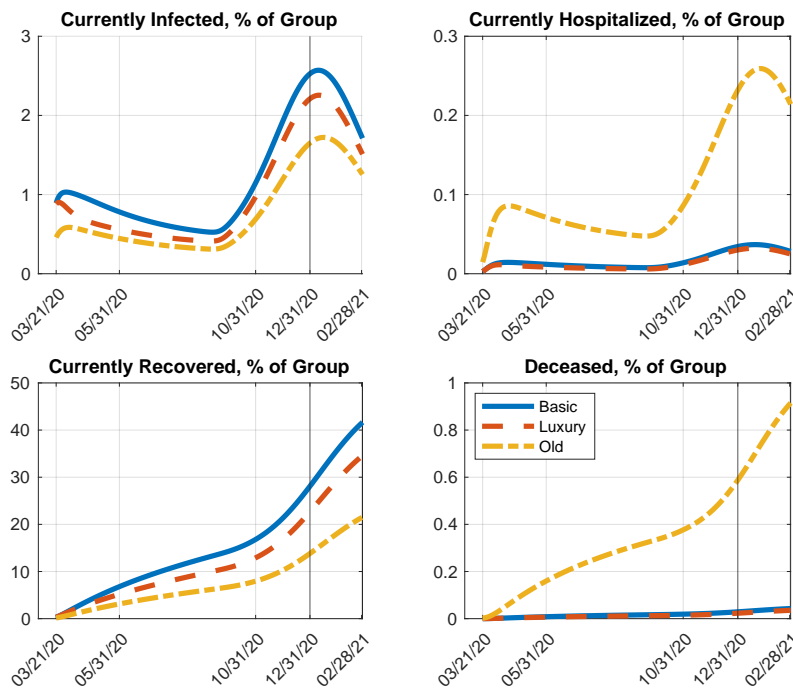


Figure 2: Share of Each Group Infected, Hospitalized, Recovered, Deceased.

One can clearly see how the young, and especially those working in the basic sector (and thus

taking care of the sick in hospital), are more likely to be infected. But a much larger share of the elderly population is hospitalized (around 0.2% of the population 65 and older at the peak of the winter 2020 wave). This differential hospitalization contributes to a massive difference in the incidence of death from COVID-19: according to our model, by February 2021 about 0.8% of the elderly population has died from the pandemic, whereas among the young the corresponding number is only 0.035%. 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 United States, according to which 80 percent of COVID-19 deaths in 2020 were among people 65 years and older.³¹

The bottom left panel shows the total share of the population (within each group) that has recovered by the end of 2020 and into 2021. By December 31, 2020, around 28% of basic workers, 23% of luxury workers and 14% of the elderly have recovered, and in the aggregate, slightly less than 24% 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 US of which we are aware.³²

5.2 The Economic Crisis through the Lens of the Model

We now turn from the health to the economic dimension of the pandemic. Figure 3 characterizes the economic recession under the baseline shutdown policy. When the government shuts down 40% of the luxury sector, it simultaneously increases the income tax rate to finance transfers to the newly unemployed: the tax rate on impact increases from 30% to 36%.

The economy plunges into a deep recession in March 2020, mostly because mitigated individuals are not working. This decline is amplified by a rise in the share of people off work because they are sick or hospitalized and by a decline in hours by those who are working but

³¹See CDC weekly updates: https://www.cdc.gov/nchs/nvss/vsrr/covid_weekly/index.htm.

³²The Institute for Health Metrics and Evaluation (IHME) estimates 18% percent of the US 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>).

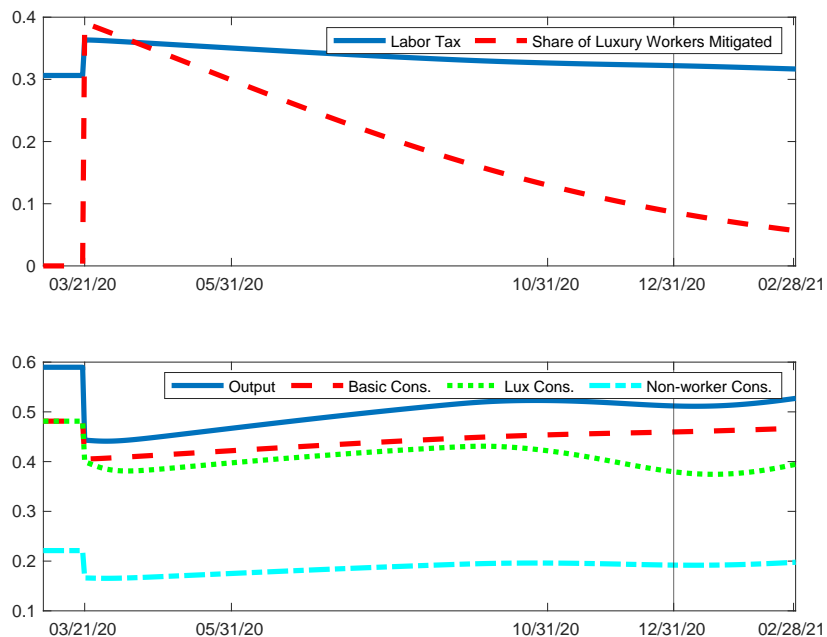


Figure 3: Mitigation, Tax Rate, Output and Consumption in Baseline Mitigation Scenario.

now face higher tax rates. Output falls by about 30% (see the bottom panel of Figure 3), and consumption of all groups tracks the path of output, with little heterogeneity in the percentage change in consumption at the onset of the crisis.

As time passes, mitigation is gradually relaxed, and output slowly recovers. However, when the fall/winter wave hits, the recovery stalls again. This reflects a mix of more sick people not working plus a decline in demand for luxury, social-contact-intensive consumption, which depresses wages and thus labor supply for workers in the luxury sector.

5.3 Consequences of Alternative Mitigation Paths

To characterize the trade-offs implied by different mitigation policies, in Figure 4 we contrast the evolution of the health pandemic as well as taxes and transfers under the benchmark policy versus a policy of no mitigation at all. The mitigation paths (the percentage of the luxury sector shuttered) are depicted in the upper left panel of Figure 4. On the health side, we see

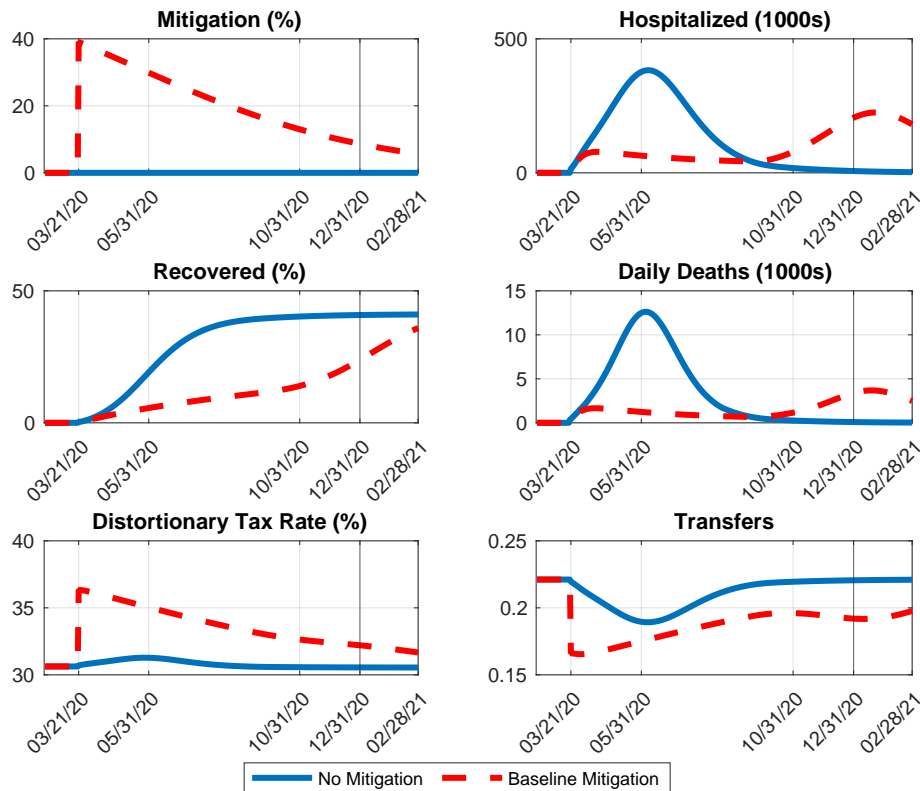


Figure 4: Outcomes: Baseline Mitigation versus Zero Mitigation.

that a policy of doing nothing (solid blue line) leads to rapid herd immunity by the fall of 2020 (see the middle left panel), at the expense of an explosion in hospitalizations and deaths in the spring and early summer of 2020. At the peak, this *laissez-faire* policy would have led to 382,000 individuals hospitalized at the beginning of June 2020 (right top panel), far outstripping the supply of emergency hospital beds available in the US. The number of deaths would have peaked not long after, at over 12,600 *daily* deaths. These additional deaths (relative to the benchmark) largely reflect the old getting sick and dying at a higher rate.

The lower two panels show the other side of the trade-off. Since more people in the luxury sector are working under the *laissez-faire* policy, taxes are lower and transfers higher than with more pronounced mitigation. However, even absent mitigation, equilibrium transfers to non-workers fall when infections surge, and the value of output is depressed by a mix of sick workers

staying home and the demand for luxury goods falling.

This trade-off between stemming the health pandemic and the economic costs of mitigation raises the question what an optimal mitigation policy looks like. In Figure 5 we characterize this optimal policy for a utilitarian government and contrast its outcomes to those under the baseline mitigation path.³³ Optimal shutdowns are milder early in the pandemic. In addition, the optimal path for mitigation inherits some of the seasonality in infection-generating rates: optimal mitigation is relaxed slightly in the summer and then ramps up again in the fall.

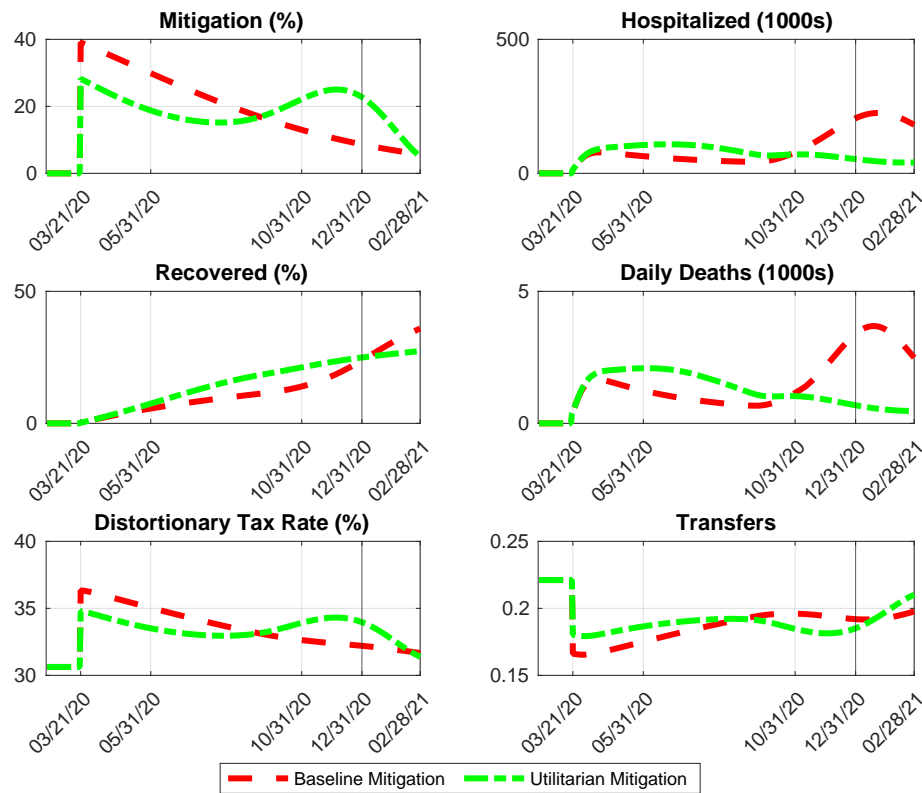


Figure 5: Outcomes: Baseline Mitigation versus Utilitarian Optimal Mitigation.

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.

³³See the Appendix for details of how we solve for the optimal policy numerically.

Table 4: Cumulative Deaths, Alternative Mitigation Paths

	No Mitigation	Baseline	Utilitarian
Through 12/31/20	1,093,705	364,750	420,854
Through 12/31/21	1,098,134	657,487	501,807

Table 4 reports total deaths through the end of 2020 and the end of 2021 – by which time the pandemic is over – under the baseline mitigation path and the utilitarian optimal one. Relative to the baseline, had the US instead followed the optimal mitigation path, the model predicts we would suffered 56,000 *more* COVID deaths in 2020, perhaps suggesting lockdowns in the first wave were too tight. But the optimal path delivers 156,000 *fewer* deaths in total, as tighter restrictions in late 2020 and early 2021 would have largely eliminated the winter wave.³⁴

Note also that because the optimal path for mitigation is flatter overall relative to the baseline, so is the share of the population working. The time path of transfers and that of tax rates display the same overall shape.

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 6 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 60% of the luxury sector and lasting well into the winter of 2020. In contrast, the young would like to see only around 10% of the luxury sector shuttered, and they would like to see lockdowns end altogether before the end of 2020. The policy chosen

³⁴Recall that subsectors in our model economy vary in terms of infection-generating contact rates. Our baseline assumption is that mitigation is directed toward the riskiest subsectors. We have also explored an alternative model for mitigation according to which subsectors are shutdown at random. In this undirected mitigation model, the terms $\alpha_w(1 - m_t)$ and $\alpha_c(1 - m_t)$ in equations (8) and (9) 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 planner chooses more mitigation when mitigation is a less effective tool. The reason is that when mitigation is less effective, infections and hospitalizations tend to run higher. Absent higher mitigation, hospital capacity would be drastically exceeded, leading to even higher mortality.

by the utilitarian government is a compromise between these divergent preferences.³⁵

Comparing the policies preferred by young workers in the basic versus the luxury sectors, the luxury workers tend to favor slightly more extensive lockdowns. At first sight this might seem surprising, because luxury workers are the ones who risk losing their jobs from mitigation, thereby suffering a loss of income and consumption. However, these workers rationally forecast that if the pandemic intensifies and hospitalizations rise, this depresses demand for luxury consumption, in turn lowering wages in the luxury sector.

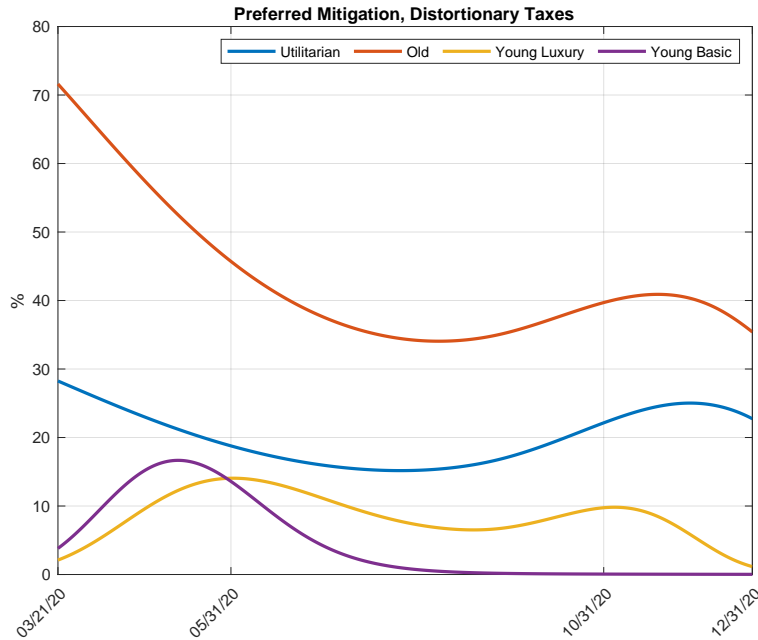


Figure 6: Preferred Mitigation Paths for Different Types.

³⁵There 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 shows up in welfare via the preference shifters \hat{u}^f and \hat{u}^e . We have solved for 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 were to emerge that would eliminate mortality, it would be optimal to set mitigation to zero.

Table 5 reports welfare gains for different types (the rows) under different mitigation paths (the columns). 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 one of the four mitigation paths plotted in Figure 6?³⁶ The table reinforces the main 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.³⁷ The young worry much 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% of consumption in the counter-factual scenario with zero mitigation. Conversely, the old see welfare gains close to 4% of consumption under their preferred extensive and persistent mitigation policy.

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

³⁶Let stars denote allocations under an alternative policy. The welfare gain for the young from switching to the alternative policy from the baseline policy is the value for ω that solves

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

Note that the equilibrium allocations in our economy have the property that

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

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.

³⁷The only reason the old do not favor 100% mitigation is the negative impact on transfers from such a policy.

Table 5: Welfare Gains (+) or Losses (-) From Mitigation, Distortionary Taxes

Group	Policy				
	None	Old	Luxury	Basic	Utilitarian
Young Basic	0.11%	-0.35%	0.20%	0.22%	0.14%
Young Luxury	0.20%	-0.79%	0.33%	0.30%	0.23%
Old	-6.18%	3.91%	-0.57%	-1.23%	2.68%
Utilitarian	-0.38%	-0.19%	0.19%	0.13%	0.39%

5.4.1 Age-Based Policies

The strong age heterogeneity in the preferences over mitigation raises the question of whether age-based policies could yield welfare gains (see also Acemoglu et al. 2021 and Brotherhood et al. 2020 for an extensive discussion of such policies). One such widely discussed policy was to isolate older people at home. In our model, the old do not work, so the only way to lower their risk of infection outside the home is to reduce their market consumption activity. We therefore consider an intervention to protect the old from infection while shopping, which is to ask the young to shop on their behalf.³⁸ To model this we remove the consumption channel as a possible source of infection for the old, and simultaneously increase the importance of the consumption channel for the young, in such a way that the total number of infections through consumption would remain unchanged, given the same distribution across health states for young and old.

Table 6: Welfare Gains (+) or Losses (-): Utilitarian Government

	Baseline	Elderly Shut In
Young Basic	0.14%	0.15%
Young Luxury	0.23%	0.29%
Old	2.68%	4.62%
Utilitarian	0.39%	0.57%

This intervention significantly benefits the old. But, once the old stop shopping, the government reduces the path for mitigation, because mitigation no longer directly protects the old

³⁸This policy can also be interpreted as introducing special shopping hours for the elderly where infection risk is minimized, in turn reducing shopping hours for the young, with more crowded and infectious stores for them.

from infection. This less drastic mitigation in turn translates into more infections. Since these infections are more heavily tilted toward the young, total deaths are reduced, however. Fewer deaths plus higher average consumption is an attractive package. Column (2) of Table 6 documents welfare gains for all groups and shows that whereas the young are essentially unaffected by a joint policy of isolation of the old plus optimal mitigation (relative to utilitarian optimal policy absent isolation), the old experience significantly larger welfare gains.

5.5 The Interaction between Mitigation and Redistribution

Thus far we have argued that there is strong divergence in policy preferences across the different population groups, since they trade off the health benefits of mitigation versus the economic costs very differently. The second main theme of this paper is that mitigation and redistribution policies interact. Concretely, the government's appetite for mitigation will be dampened when it is harder to adjust redistribution policy so as to soften the impact of shutdowns on inequality.

To make this point, we now compare optimal policies under the benchmark fiscal constitution to three alternatives, starting with the lump-sum transfer system studied theoretically in Proposition 2. Figure 7 shows the preferred mitigation paths for all three population groups as well as the utilitarian optimum, under lump-sum group-specific transfers (red dashed line) as well as under the benchmark fiscal system with distortionary taxes (blue solid line). The figure also shows optimal mitigation under two alternative versions of the distortionary tax economy: one in which tax rates are held fixed over time (green dotted line) and a second in which transfers are held fixed (black dash-dotted line).³⁹

If the government has access to group-specific lump-sum taxes and transfers, then taxation does not distort labor supply choices, redistribution is costless, and the utilitarian government optimally equates the marginal utility of consumption across all agents at each date. With our

³⁹One could argue that the most relevant fiscal policy scenario is one where the additional transfers due to mitigated workers are financed through government debt, which is serviced through higher taxes on the current young and future generations. With lump-sum taxes the timing of these taxes would not matter (absent intergenerational redistribution), owing to Ricardian equivalence, so for this economy, abstracting from debt is not an issue. With distortionary taxes, debt would allow the planner could to smooth higher marginal tax rates over time (Lucas and Stokey, 1983).

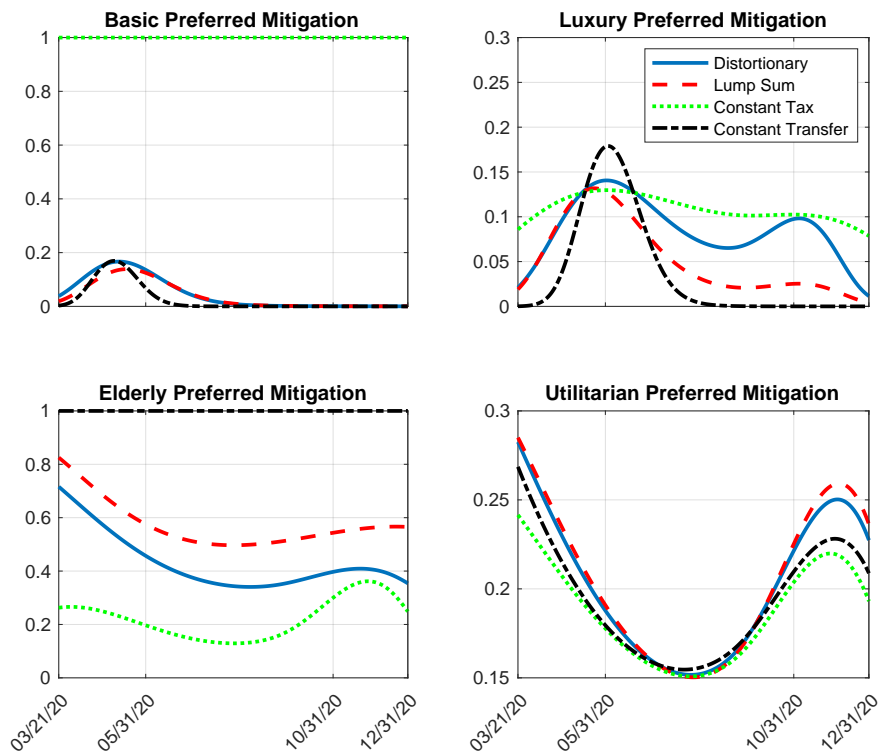


Figure 7: Preferred Mitigation: Alternative Fiscal Rules

utility function, this requires equating the *level* of utility from consumption and labor supply. Any difference in preferred policies across types under lump-sum taxes then reflects exclusively differential concerns about the health consequences of COVID-19.

The red dashed lines in Figure 7 demonstrate that the old care about these health consequences much more than the young and consequently prefer more extensive mitigation.⁴⁰ Crucially, comparing the lump-sum tax system to the distortionary benchmark shows that the utilitarian optimal path calls for more mitigation (on average) when redistribution via taxation does not distort labor supply. This was the content of Proposition 2.⁴¹ The elderly are especially

⁴⁰For health reasons, young basic workers, who care for the sick and thus have higher infection risk, prefer slightly more mitigation than young luxury workers, although this effect is quantitatively small and visible only in the more severe second wave of the pandemic.

⁴¹Note 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.

keen on extra mitigation with lump-sum taxation, because now increasing taxes and transfers does not increase distortions (in contrast to the baseline economy), which implies a smaller elasticity of transfers to mitigation. The difference in preferred mitigation under the lump-sum and distortionary systems is less pronounced for the young. Young luxury workers support more mitigation under the benchmark tax system (especially in the second wave) as a tool to protect demand for the social-contact-intensive luxury goods they produce. With lump-sum transfers, period utility is equalized across groups, and thus luxury workers are already effectively insured against luxury price declines in the winter 2020 COVID-19 wave.

The second point we want to make with Figure 7 is that although the *utilitarian* optimal mitigation path is fairly robust to differences in the tax-transfer system, group-specific mitigation policy preferences vary strongly with the fiscal regime, since these heterogeneous groups care about fundamentally different aspects of fiscal policy. To this end, consider the two additional fiscal constitutions plotted that hold fixed, respectively, the tax rate and the level of transfers.⁴²

As discussed before, the old care about fiscal policy only to the extent that it impacts their transfers, and if these are guaranteed to be constant, then they support 100% mitigation for health reasons (lower left panel). At the other extreme, workers in the basic sector receive transfers only if they fall sick, but they face potentially much higher taxes when transfers are held constant and thus support very little mitigation in that scenario (upper-left panel). In contrast, with constant taxes this group would happily support 100% mitigation, while the implied lower transfers curb even the old's enthusiasm for significant shutdowns (see the dotted green lines in the upper-left and lower left panels). Workers in the luxury sector are in the middle of these extremes, as they are both tax payers (if their businesses are not shut down) and transfer recipients (if they are mitigated), and thus their policy preferences tend to be in between those of the old and the young basic workers.

⁴²The first scenario is of policy interest when it is difficult to rapidly increase taxes and transfers because it takes time to change fiscal policy or when it is difficult because the government has limited fiscal space. Given a constant tax rate and a balanced budget requirement, increasing mitigation will then dictate a large decline in transfers. The second scenario is relevant if it is politically difficult to reduce public transfers to retirees or to cut benefits to the unemployed. Now increasing mitigation will dictate a large rise in the equilibrium tax rate.

One especially interesting comparison is between the lump-sum tax economy and the fixed tax rate economy. In both cases, changing the path for mitigation does not affect the magnitude of distortions to labor supply or output. But in the lump-sum tax economy changing the path for mitigation also does not change equilibrium inequality in flow utility from consumption and labor supply – the government always eliminates any such inequality – while in the fixed tax rate economy, inequality in consumption and flow utility widens as mitigation and the fraction of non-workers rise. The lower-right panel of Figure 7 illustrates that this concern about rising inequality substantially tempers the utilitarian government’s appetite for mitigation.

Taken together these comparisons clearly illustrate that mitigation and redistribution are complementary policies. When redistribution is easier, stronger mitigation is optimal. However, the differences between the optimal utilitarian paths for mitigation under different fiscal rules are quantitatively relatively modest. This result is reassuring, in that it suggests that utilitarian optimal mitigation policy is quite robust to details of tax and transfer system, even though the preferences of different household types vary strongly across different systems. At the same time our exploration suggests that the extent of political support for mitigation is likely to be highly sensitive to the fiscal constitution. Given the baseline tax system, only the elderly support strong mitigation. Since they are in a minority in our economy (15 percent of the population), the median voter will be young and would choose less mitigation than the utilitarian government. One way to build political support for more mitigation would be to promise to hold taxes fixed (or close to fixed), thereby incentivizing the young to support stronger mitigation.

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

Having characterized the desired mitigation path for 2020, we now study optimal mitigation and redistribution policy in a pandemic when a vaccine becomes available. By the beginning of 2021, it was clear that a roll-out of effective vaccines that would eventually end the pandemic was imminent. We now ask how the emergence and diffusion of a vaccine that protects individuals both from illness as well as from spreading the virus to others impacts the optimal lockdown policy and health as well as economic outcomes. To do so, we assume that the vaccine is perfectly

effective and is rolled out among the different age groups in the model in accordance with the actual pattern for the US in the first half of 2021.⁴³ When susceptible people are vaccinated, by assumption they immediately transition to the recovered state in the epidemiological model, thus pushing the economy closer to herd immunity.⁴⁴

We make one additional change for our 2021 simulations relative to the backward-looking simulations 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 vaccinations proceed, people become less concerned about contracting and transmitting the virus and thus return to pre-COVID social behavior.⁴⁵ We set the constant of proportionality so that if everyone were to be vaccinated – implying a fully recovered population – the basic reproduction number R_0 would return to its initial, pre-March 20, 2020 value. See Figure 13 in the Appendix for the model-implied evolution of the basic reproduction number.

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. Equipped with our initial condition and the deterministic time paths for vaccinations by age, we explore alternative mitigation paths in the presence of a vaccine roll-out.

Figure 8 displays preferred mitigation paths in the presence of a vaccine roll-out in the first half of 2021. It is the counterpart to Figure 6 for 2020. Comparing these two figures, we derive two basic insights. First, the ranking of policy preferences remains unchanged in

⁴³Model 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 been vaccinated.

⁴⁴Very recent evidence suggests that fully vaccinated individuals can still transmit the virus, especially the “delta” variant, though they are much less likely to do so than infected but unvaccinated individuals.

⁴⁵We acknowledge that this impact of changes in individual behaviors in response to the evolution of the pandemic is not fully micro-founded. Farboodi et al. (2021) and Engle et al. (2021) are important examples of models that endogenize this behavioral response.

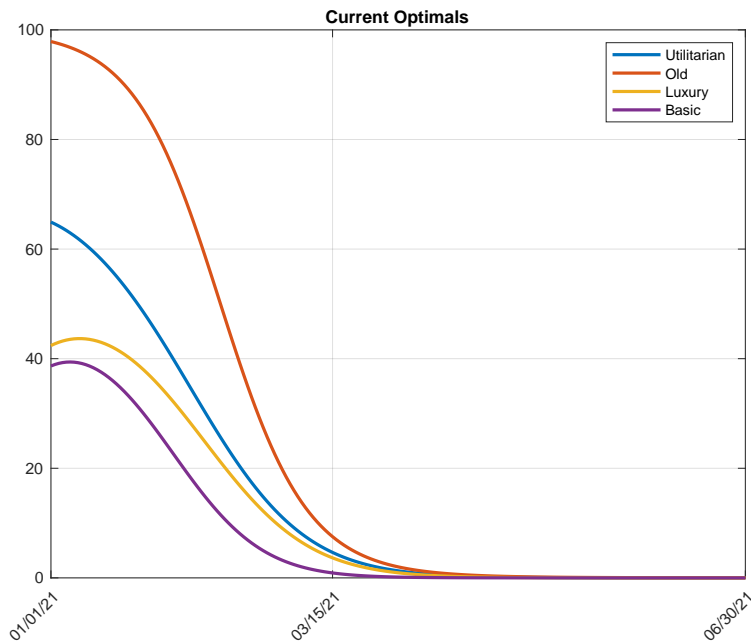


Figure 8: Preferred Mitigation in the Presence of a Vaccine

the presence of a vaccine, with the old preferring more mitigation than the young working in both sectors and basic workers supporting a quicker reduction of lockdowns as the winter wave of infection subsides. As before, the utilitarian optimum lies in between these group-specific policy preferences. Second, and more importantly, the optimal *level* of mitigation is significantly higher in early 2021 than during 2020, and the reopening proceeds at a more rapid pace. The anticipated diffusion of a vaccine means that infections prevented by mitigation in early 2021 will never occur. Absent a vaccine, in contrast, mitigation primarily delays infections rather than eliminating them altogether.

That the arrival of the vaccine is indeed responsible for the stronger preference for mitigation is demonstrated by Figure 9, which shows the optimal mitigation paths for 2021 in the presence and absence of vaccine deployment, in addition to the benchmark path for mitigation. Throughout January optimal mitigation with a vaccine is about 15 percentage points higher than without, but then mitigation falls much more quickly in February through May 2021, as rising vaccination rates translate into rapidly declining infections. Thus, news about a successful

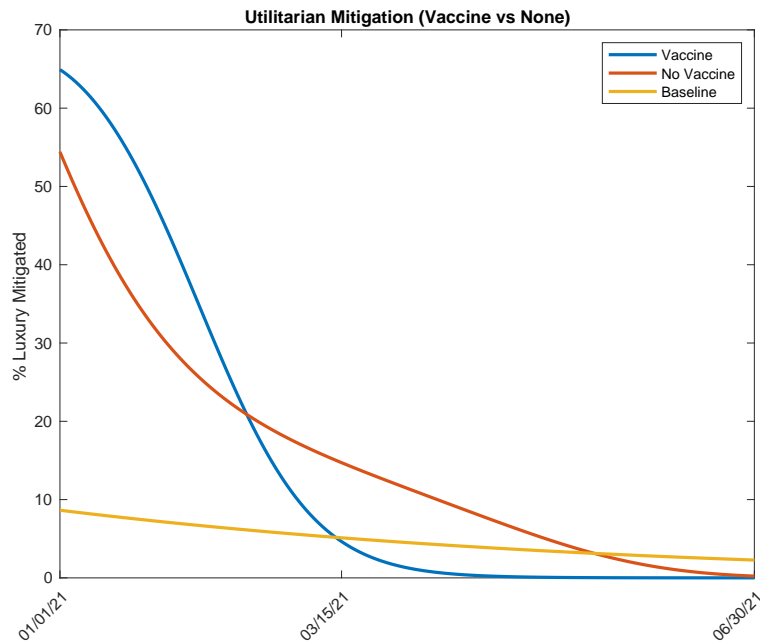


Figure 9: Preferred Mitigation with Vaccine: Difference Compared With No Vaccine

future vaccine roll out changes optimal mitigation policy quite dramatically relative to the no vaccine scenario.⁴⁶

We now compare health and economic outcomes under three alternative scenarios for the deployment of vaccines and the path for mitigation. In the first (yellow line), our approximation of actual events, vaccines are gradually rolled out, and the path for mitigation is the baseline one plotted in Figure 9. In the second scenario there are no vaccines, and the path for mitigation is the optimal one for the no-vaccine scenario (the red line in Figure 9). In the third scenario, vaccines are rolled out, and the path for mitigation is the optimal one conditional on vaccine deployment (the blue line depicted in Figure 8).

⁴⁶Note that relative to the spring and summer of 2020, optimal (utilitarian) mitigation is higher in early 2021 even in the absence of a vaccine, by about 15-20 percentage points. This is due to the fact that the second wave in the winter of 2020/21 is more massive, with more infected individuals in the population and the hospitals at risk of being overrun, as well as the assumption that infection rates respond to the state of the pandemic in 2021, as described at the beginning of this section.

6.1 Health Dynamics in the Presence of a Vaccine Roll-Out

Figure 10 plots daily deaths under the three scenarios and the cumulative share of the population that has recovered from COVID (either via infection or via vaccination). It shows that in the presence of a vaccine, daily deaths rapidly converge to zero by the middle of 2021 as a large share of the population becomes vaccinated or has recovered from the disease. The bottom panel shows, however, that the mitigation path chosen is a strong determinant of the number of casualties on the way to this eventual outcome, with the optimal mitigation path (which implies much more significant lockdowns in early 2021) avoiding about 175,000 deaths relative to the baseline scenario. Interestingly, even in the absence of a vaccine the optimal mitigation (which is substantial and lasts well into the summer of 2021) delivers health outcomes that are not too much worse than those with a vaccine (compare the red and blue lines in Figure 10), but at the expense of longer, economically costly lockdowns.

6.2 The Evolution of the Economy with Vaccination

The economic consequences of the availability of vaccines and the optimal mitigation policy associated with it are depicted in Figure 11, which shows the optimal tax rate (upper panel) and the deviation of output from its value at the end of 2020 level (lower panel) under the same three scenarios described above. Under the two scenarios in which the mitigation path is chosen optimally, lockdowns are initially much higher than under the baseline path and are associated with higher tax rates and depressed output at the start of 2021.

As the economy moves into spring, the economy recovers fast in the scenarios with vaccines, because fewer workers are sick and because demand in the luxury sector recovers as hospitalizations fall. When mitigation is chosen optimally, it is rapidly relaxed as the share of the population vaccinated and protected from infection rises. This further fuels the recovery, both directly, as previously mitigated workers return to work, and indirectly, as lower tax rates encourage additional labor supply. In contrast, in the scenario without a vaccine, the recovery in 2021 proceeds much more gradually. Thus, from the perspective of our model, getting the

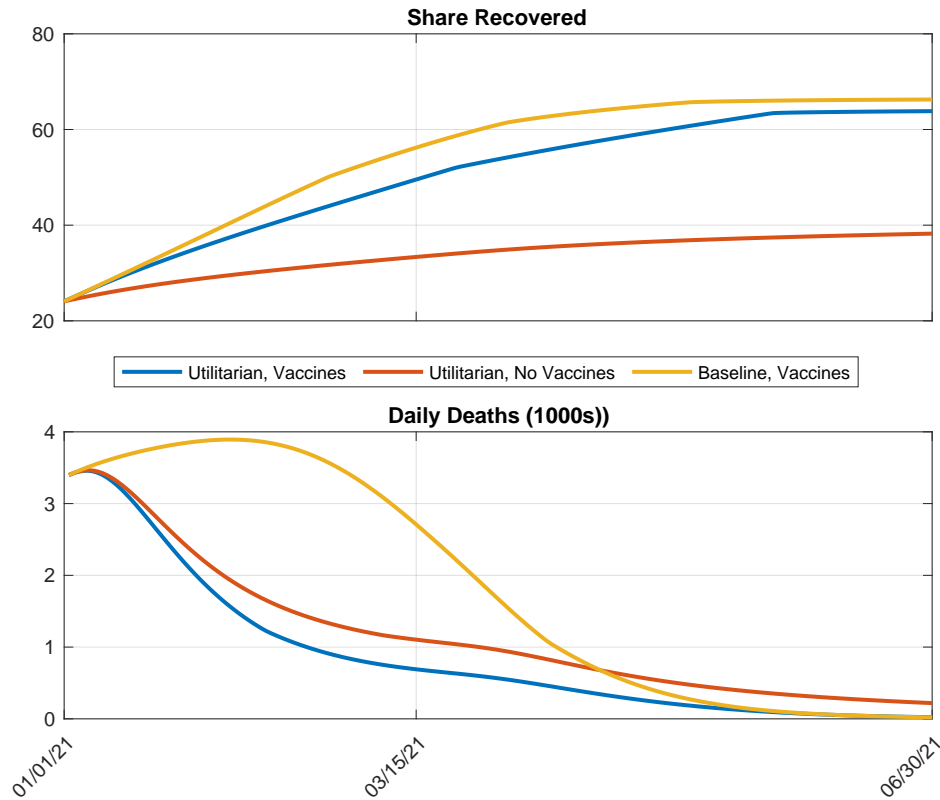


Figure 10: Evolution of Recovered Population Share, Daily Deaths in 2021

pandemic under control with a vaccine is the key to a rapid economic recovery.

Table 7 summarizes the welfare gains from the introduction of the vaccine. The different columns of the table correspond to the three scenarios plotted in Figures 10 and 11. The rows report four outcomes, each relative to the benchmark scenario in which no vaccines ever arrive and mitigation follows the baseline path. These outcomes 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, and the increase in total GDP for 2021.⁴⁷

The first column indicates that welfare gains from the vaccine alone are large and equivalent

⁴⁷Welfare gains are calculated as the permanent percent increase in the utility kernel (which is linear in consumption) needed to make the utilitarian government (or the old) indifferent between a given scenario and the no vaccine plus baseline mitigation benchmark alternative.

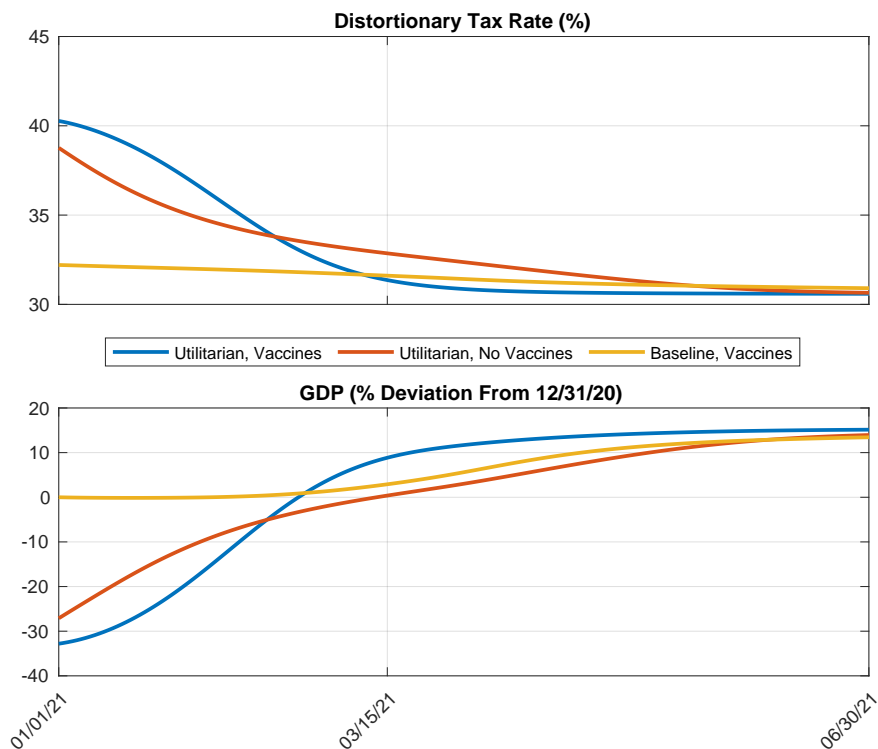


Figure 11: Evolution of Taxes and GDP (Deviation from Benchmark) in 2021

to increasing consumption permanently by one third of a percent for individuals' remaining lifetimes.⁴⁸ The gains for the old are much larger, at 3 percent of consumption. Vaccines, holding mitigation policy constant, save almost 160,000 lives and boost GDP in 2021 by 1.1 percent, owing to reduced sickness and stronger demand for luxury consumption.

The second column shows that absent a vaccine, the government can attain similar welfare gains if it moves from the baseline (very modest) mitigation path to the optimal one plotted in Figure 9. However, the source of welfare gains is quite different: a fairly extensive and persistent path for mitigation saves even more lives (260,000) but at the cost of a recession in which output falls by an additional 1.09 percent.

The third column of the table shows the gains from the most favorable scenario: a vaccine

⁴⁸Given U.S. consumption of around \$15Tr, this corresponds to \$50bn per year on a flow basis. For comparison, the (one time) budget for Operation Warp speed was \$10bn.

Table 7: Welfare Gains From Vaccine Introduction

	Vaccine Baseline Mitigation	No Vaccine Optimal Mitigation	Vaccine Optimal Mitigation
Utilitarian Welfare	0.34%	0.35%	0.64%
Old Welfare	2.95%	3.82%	5.91%
Deaths Avoided	159,583	260,430	335,123
GDP Gain, 2021	1.10%	-1.09%	0.12%

is deployed, and the path for mitigation is chosen optimally. Now, relative to the vaccine-only scenario, welfare gains double, and more than twice as many lives are saved, while output per capita is mildly higher than in the benchmark scenario. We conclude that while having a vaccine is (unsurprisingly) hugely beneficial for welfare, even with a vaccine, it remains very important to choose an appropriate path for economic mitigation. That path (Figure 9) involves a hard lockdown to save lives in the short run, coupled with rapid relaxation to save the economy as vaccinations progress and infections slow (compare the last row across columns 2 and 3) in Table 7.

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 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 relatively mild, at the cost of more adverse health outcomes. Our baseline calibration suggests that the shutdown actually in place in the spring of 2020 was too extensive, but it was also lifted too quickly, relative to what a utilitarian government would choose. 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 of 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 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, however, 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 future pandemic health crises 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. While we think of this as an important area of future research, it is our hope that these conjectures are not put to an empirical test anytime soon again.

References

- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. WHINSTON (2021): "Optimal Targeted Lockdowns in a Multi-Group SIR Model," American Economic Review: Insights, forthcoming.
- ARGENTE, D., F. ALVAREZ, AND F. LIPPI (2021): "A Simple Planning Problem for COVID-19 Lockdown, Testing, and Tracing," American Economic Review: Insights, forthcoming.
- 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).
- BERGER, D., K. HERKENHOFF, C. HUANG, AND S. MONGEY (2020): "Testing and Re-opening in an SEIR model," Review of Economic Dynamics, forthcoming.
- 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 (2020): "Integrated Epi-Econ Assessment," Working Paper 28282, National Bureau of Economic Research.
- 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-cenarios-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, forthcoming.

- EICHENBAUM, M., S. REBELO, AND M. TRABANDT (2020): "The Macroeconomics of Epidemics," National Bureau of Economic Research Working Paper 26882.
- 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*, forthcoming.
- FERNANDEZ-VILLAVARDE, J. AND C. I. JONES (2020): "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities," National Bureau of Economic Research Working Paper 27128.
- GARRIGA, C., R. E. MANUELLI, AND S. SIDDHARTHA (2021): "Optimal Management of an Epidemic: Lockdown, Vaccine and Value of Life," Working Paper 046-A, Federal Reserve Bank of St. Louis Working Paper Series.
- GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RÍOS-RULL (2020): "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," London, Centre for Economic Policy Research, DP 15980.
- 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 (2020): "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" Working Paper 26918, National Bureau of Economic Research.
- 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.
- HONG, J. H., J. PIJOAN-MAS, AND J.-V. RIOS-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 (2020): "Macroeconomic Dynamics and Reallocation in an Epidemic," Centre for Economic Policy Research Discussion Paper 14607.

- 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," Unpublished 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>.
- MULLIGAN, C. B. (2020): "Deaths of Despair and the Incidence of Excess Mortality in 2020," National Bureau of Economic Research Working Paper 28303.
- 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," Unpublished Manuscript, 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.
- XU, J., G. WATSON, D. XIONG, M. A. SUCHARD, AND C. M. RAMIREZ (2020): "COVID-19 Infection and Mortality across Workplaces and Intervention Effects," Mimeo, UCLA.

Appendix: Not for Publication

A Details of the SAFER Model

This section summarizes the details of the remainder of the epidemiological block. It is fairly mechanical and simply describes the transition of individuals through the health states (asymptomatic, fever-suffering, hospitalized, and recovered) once they have been infected. Equations (33) to (35) describe the change in the measure of asymptomatic individuals. There is entry into that state from the newly infected flowing in from the susceptible state (as described above). Exit from this state to developing a fever occurs at rate σ^{yaf} (σ^{oaf}) for the young (old), and exit to the recovered state occurs at rate σ^{yar} (σ^{oar}) for the young (old). Note that someone who recovers at this stage will never know that she contracted the virus.

For individuals suffering from a fever, equations (36) to (38) show that for the young there is entry from the asymptomatic state and exit to the hospitalized state at rate σ^{yfe} , and to the recovered state at rate σ^{yfr} , with analogous expressions for the old. Equations (39) to (41) describe the movements of those in emergency care, showing entry from those with a fever and exits to death and recovery. The death rate is $\sigma^{yed} + \varphi$, while the recovery rate is $\sigma^{yer} - \varphi$, where φ , described in equation 13 in the main text, is a term related to hospital overuse. Equations (42) to (44) display the evolution of the measure of the recovered population, which features only entry and is an absorbing state. So is death, with the evolution of the deceased population being determined by $\dot{x}^{ybd} = (\sigma^{yed} + \varphi)x^{ybe}$, $\dot{x}^{yld} = (\sigma^{yed} + \varphi)x^{yle}$, and $\dot{x}^{od} = (\sigma^{oed} + \varphi)x^{oe}$. We record them separately from the recovered (who work), since they play no further role in the model.

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

$$\dot{x}^{yba} = -\dot{x}^{ybs} - (\sigma^{yaf} + \sigma^{yar}) x^{yba} \quad (33)$$

$$\dot{x}^{y\ell a} = -\dot{x}^{y\ell s} - (\sigma^{yaf} + \sigma^{yar}) x^{y\ell a} \quad (34)$$

$$\dot{x}^{oa} = -\dot{x}^{os} - (\sigma^{oaf} + \sigma^{oar}) x^{oa} \quad (35)$$

$$\dot{x}^{ybf} = \sigma^{yaf} x^{yba} - (\sigma^{yfe} + \sigma^{yfr}) x^{ybf} \quad (36)$$

$$\dot{x}^{y\ell f} = \sigma^{yaf} x^{y\ell a} - (\sigma^{yfe} + \sigma^{yfr}) x^{y\ell f} \quad (37)$$

$$\dot{x}^{of} = \sigma^{oaf} x^{oa} - (\sigma^{ofe} + \sigma^{ofr}) x^{of} \quad (38)$$

$$\dot{x}^{ybe} = \sigma^{yfe} x^{ybf} - (\sigma^{yed} + \sigma^{yer}) x^{ybe} \quad (39)$$

$$\dot{x}^{yle} = \sigma^{yfe} x^{y\ell f} - (\sigma^{yed} + \sigma^{yer}) x^{yle} \quad (40)$$

$$\dot{x}^{oe} = \sigma^{ofe} x^{of} - (\sigma^{oed} + \sigma^{oer}) x^{oe} \quad (41)$$

$$\dot{x}^{ybr} = \sigma^{yar} x^{yba} + \sigma^{yfr} x^{ybf} + (\sigma^{yer} - \varphi)x^{ybe} \quad (42)$$

$$\dot{x}^{y\ell r} = \sigma^{yar} x^{y\ell a} + \sigma^{yfr} x^{y\ell f} + (\sigma^{yer} - \varphi) x^{y\ell e} \quad (43)$$

$$\dot{x}^{or} = \sigma^{oar} x^{oa} + \sigma^{ofr} x^{of} + (\sigma^{oer} - \varphi) x^{oe} \quad (44)$$

$$(45)$$

B Details of the Theoretical Results of Section 2

In this section we provide the detailed derivations for the theoretical results in Section 3.

B.1 Proof of Proposition 1

First, we characterize the optimal redistributive tax and transfer policy, as characterized in Proposition 1.

B.1.1 Household Decisions

As stated in the main text, the household maximization problem for those not mitigated is

$$\begin{aligned} & \log \left(c - \frac{h^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}} \right) \\ c &= c^b + c^\ell (1 - \xi_c) \\ c^b + (1 - \xi_c) c^\ell &= (1 - \tau) h w^i \end{aligned}$$

Given the relative price $p = (1 - \xi_c)$, households are indifferent between consuming the basic and the luxury good. Optimal labor supply h and total consumption c satisfy the first order conditions

$$\begin{aligned} \frac{1}{c - \frac{h^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}}} &= \lambda^i \\ \frac{h^{\frac{1}{\chi}}}{c - \frac{h^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}}} &= \lambda^i (1 - \tau) w^i \end{aligned}$$

with solution

$$\begin{aligned} h^i &= [(1 - \tau) w^i]^\chi \\ c^i &= (1 - \tau) h^i w^i = [(1 - \tau) w^i]^{1+\chi} \end{aligned}$$

for $i \in \{b, \ell\}$ and $w^i = 1$, $w^\ell = 1 - \xi_c$. Utility from this allocation is determined by

$$U^i = \log \left(c^i - \frac{(h^i)^{1+\frac{1}{\chi}}}{1+\frac{1}{\chi}} \right) = \log \left([(1 - \tau) w^i]^{1+\chi} - \frac{[(1 - \tau) w^i]^{1+\chi}}{1+\frac{1}{\chi}} \right)$$

$$= \log \left(\frac{[(1-\tau)w^i]^{1+\chi}}{1+\chi} \right) = \log \left(\frac{1}{1+\chi} \right) + (1+\chi) \log(1-\tau) + (1+\chi) \log(w^i)$$

and therefore, for each group

$$U^b(\tau) = \log \left(\frac{[1-\tau]^{1+\chi}}{1+\chi} \right) = \log \left(\frac{1}{1+\chi} \right) + (1+\chi) \log(1-\tau)$$

$$U^\ell(\tau) = \log \left(\frac{[(1-\tau)(1-\xi_c)]^{1+\chi}}{1+\chi} \right) = \log \left(\frac{1}{1+\chi} \right) + (1+\chi) \log(1-\tau) + (1+\chi) \log(1-\xi_c)$$

For non-working households, as a direct consequence of their budget constraint we have

$$\begin{aligned} c &= c^b + c^l(1-\xi_c) = T \\ U^n &= \log(c) = \log(T). \end{aligned}$$

Using the government budget constraint we can write transfers as a function of the tax rate τ as follows:

$$\begin{aligned} T &= \frac{(\tau-g)[Y^b + pY^\ell]}{x^n(m)} = \frac{(\tau-g)[x^{bw}[(1-\tau)^\chi + x^{\ell w}(m)(1-\xi_c)^{1+\chi}][1-\tau]^\chi]}{x^n(m)} \\ &= (\tau-g)(1-\tau)^\chi \frac{[x^{bw} + x^{\ell w}(m)(1-\xi_c)^{1+\chi}]}{x^n(m)} = \frac{(\tau-g)(1-\tau)^\chi}{x^n(m)} [x^{bw} + x^{\ell w}(m)(1-\xi_c)^{1+\chi}] \\ &= \frac{(\tau-g)(1-\tau)^\chi}{x^n(m)} Y^{LS}(m) \end{aligned}$$

where

$$Y^{LS}(m) = [x^{bw} + x^{\ell w}(m)(1-\xi_c)^{1+\chi}] = (1-\tau)^{-\chi} [Y^b + pY^\ell] = (1-\tau)^{-\chi} [Y(m, \tau)]$$

is aggregate GDP under lump-sum taxes. Thus we can write

$$U^n(\tau, m) = \log(T) = \log(\tau-g) + \chi \log(1-\tau) + \log \left[\frac{Y^{LS}(m)}{x^n(m)} \right]$$

as in the main text.

B.1.2 The Static Optimization Problem

The objective of the government is to maximize, by choice of τ , static social welfare

$$\begin{aligned} W(\tau, m) &= x^{bw} U^b + x^{\ell w}(m) U^\ell + x^n(m) U^n = x^{bw} \left[\log \left(\frac{1}{1+\chi} \right) + (1+\chi) \log(1-\tau) \right] \\ &\quad + x^{\ell w}(m) \left[\log \left(\frac{1}{1+\chi} \right) + (1+\chi) \log(1-\tau) + (1+\chi) \log(1-\xi_c) \right] \\ &\quad + x^n(m) \left[\log(\tau-g) + \chi \log(1-\tau) + \log \left[\frac{Y^{LS}(m)}{x^n(m)} \right] \right] \end{aligned}$$

$$\begin{aligned}
&= \left((1+\chi) [x^{bw} + x^{\ell w}(m)] + \chi x^n(m) \right) \log(1-\tau) + x^n(m) \log(\tau-g) \\
&\quad + [x^{bw} + x^{\ell w}(m)] \left[\log\left(\frac{1}{1+\chi}\right) \right] + x^{\ell w}(m)(1+\chi) \log(1-\xi_c) + x^n(m) \log\left[\frac{Y^{LS}(m)}{x^n(m)}\right] \\
&= (x^w(m) + \chi x) \log(1-\tau) + x^n(m) \log(\tau-g) \\
&\quad + x^w(m) \left[\log\left(\frac{1}{1+\chi}\right) \right] + x^{\ell w}(m)(1+\chi) \log(1-\xi_c) + x^n(m) \log\left[\frac{Y^{LS}(m)}{x^n(m)}\right] \\
&= \Gamma(\tau, m) + \Psi(m)
\end{aligned} \tag{46}$$

defined in equation (23) of the main text, where

$$\begin{aligned}
\Gamma(\tau, m) &= (x^w(m) + \chi x) \log(1-\tau) + x^n(m) \log(\tau-g) \\
\Psi(m) &= x^w(m) \left[\log\left(\frac{1}{1+\chi}\right) \right] + x^{\ell w}(m)(1+\chi) \log(1-\xi_c) + x^n(m) \log\left[\frac{Y^{LS}(m)}{x^n(m)}\right]
\end{aligned}$$

$\Psi(m)$ is a function solely of m and thus irrelevant for maximization with respect to τ . The first order condition with respect to τ reads as $\Gamma'(\tau, m) = 0$ and thus:

$$\begin{aligned}
\frac{x^w(m) + \chi x}{1-\tau} &= \frac{x^n(m)}{\tau-g} \\
\tau^*(m; \chi, g) &= \frac{x^n(m) + g(x^w(m) + \chi x)}{(1+\chi)x} = \frac{(1-g)\mu^n(m)}{1+\chi} + g
\end{aligned}$$

where we recall that $\mu^n(m) = x^n(m)/x$ is the share of the population that is not working. Plugging in the optimal tax rate we obtain the optimal associated transfers as:

$$T^* = \left(\frac{1-g}{1+\chi}\right)^{1+\chi} (\chi + \mu^w(m))^\chi \frac{Y^{LS}(m)}{x} \tag{47}$$

which concludes the proof of Proposition 1.

B.2 Evaluating Static Social Welfare

In order to derive the static social welfare function we now plug the optimal tax-transfer policy (τ^*, T^*) into (46). First we note that per capita output with undistorted labor supply is given by

$$\begin{aligned}
\frac{Y^{LS}(m)}{x} &= \frac{[x^{bw} + x^{\ell w}(m)(1-\xi_c)^{1+\chi}]}{x} \\
&= [(1-\xi_c)^{1+\chi}] \cdot \frac{x^{bw}(1-\xi_c)^{-(1+\chi)} + x^{\ell w}(m)}{x} \\
&= [(1-\xi_c)^{1+\chi}] \cdot \frac{x^{bw} [(1-\xi_c)^{-(1+\chi)} - 1] + x^{\ell w}(m) + x^{bw}}{x} \\
&= [(1-\xi_c)^{1+\chi}] \cdot \left(\frac{x^{bw} [(1-\xi_c)^{-(1+\chi)} - 1] + x^{\ell w}(m) + x^{bw}}{x} \right)
\end{aligned}$$

$$= [(1 - \xi_c)^{1+\chi}] \cdot \left(\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m) \right)$$

To evaluate welfare we also note that

$$1 - \tau^* = \frac{(1 + \chi - \mu^n(m))(1 - g)}{1 + \chi} = \frac{(\mu^w(m) + \chi)(1 - g)}{1 + \chi}$$

$$\tau^* - g = \frac{(1 - g)\mu^n(m)}{1 + \chi}$$

B.2.1 Static Social Welfare is Monotonically Declining in Mitigation

From these expressions we can write welfare of all three subgroups of the population, and therefore total static welfare from consumption-labor allocations as a function of mitigation m solely through the share $\mu^w(m)$ of workers in the population. This is equation 28 in the main text:

$$U^n(\mu^w(m)) = \log(\tau - g) + \chi \log(1 - \tau) + \log \left[\frac{Y^{LS}(m)}{x^n(m)} \right]$$

$$= (1 + \chi) \log \left(\frac{1 - g}{1 + \chi} \right) + \chi \log(\chi + \mu^w(m))$$

$$+ \log \left[[(1 - \xi_c)^{1+\chi}] \cdot \left(\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m) \right) \right] \quad (48)$$

$$U^b(\mu^w(m)) = (1 + \chi) \log \left(\frac{1 - g}{1 + \chi} \right) + (1 + \chi) \log(\chi + \mu^w(m)) - \log(1 + \chi) \quad (49)$$

$$U^\ell(\mu^w(m)) = U^b(\mu^w(m)) + (1 + \chi) \log(1 - \xi_c) \quad (50)$$

Since the share of workers $\mu^w(m)$ is strictly decreasing in mitigation and $Y^{LS}(m)$ is strictly decreasing in mitigation, it is immediate that *welfare of all three groups is strictly increasing in $\mu^w(m)$ and thus strictly decreasing in mitigation*. Also note that m affects welfare of all groups exclusively through the share of workers $\mu^w(m)$. However, these results are not sufficient to show that social welfare is strictly decreasing in mitigation m . For that we also need to guarantee that welfare of the n -type is strictly lower than that of the ℓ types so that a shift of population mass from workers to non-workers does not raise social welfare. We now show that under Assumption 1 in the main text $U^n(\mu^w(m)) < U^\ell(\mu^w(m))$ for all $m \in [0, 1]$.

Lemma 1 *Let Assumption 1 be satisfied. Then for all $m \in [0, 1]$ we have $U^n(\mu^w(m)) < U^\ell(\mu^w(m))$.*

Proof 1 *Using equations (48) and (50) we have*

$$U^\ell - U^n = \log(\chi + \mu^w(m)) - \log(1 + \chi) + (1 + \chi) \log(1 - \xi_c) - \log \left[\frac{Y^{LS}(m)}{x} \right]$$

$$= \log(\chi + \mu^w(m)) - \log(1 + \chi) - \log \left[\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m) \right]$$

We want to show that under Assumption 1 we have $U^\ell - U^n > 0$ for all $m \in [0, 1]$. But

$$\begin{aligned}
& U^\ell - U^n > 0 \\
\Leftrightarrow & \log\left(\frac{\chi + \mu^w(m)}{1 + \chi}\right) > \log\left[\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1\right] + \mu^w(m)\right] \\
\Leftrightarrow & \frac{\chi + \mu^w(m)}{1 + \chi} - \mu^w(m) > \mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1\right] \\
\Leftrightarrow & \frac{\chi}{1 + \chi}(1 - \mu^w(m)) > \mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1\right]
\end{aligned}$$

Since $\mu^w(m)$ is strictly decreasing in m the above inequality holds for all $m \in [0, 1]$ if it holds for $m = 0$, that is, if

$$\begin{aligned}
\Leftrightarrow & \frac{\chi}{1 + \chi}(1 - \mu^w(m = 0)) > \mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1\right] \\
\Leftrightarrow & \frac{\chi(1 - \mu^w(m = 0))}{\mu^{bw}(1 + \chi)} + 1 > (1 - \xi_c)^{-(1+\chi)}
\end{aligned}$$

that is, as long as the preference shock is sufficiently small, $\xi_c < \bar{\xi}_c$ where $\bar{\xi}_c$ is defined in Assumption 1. But this is exactly the content of Assumption 1.

This result establishes that static social welfare is strictly decreasing in mitigation since higher mitigation reduces the number of tax-paying workers, increases the (optimal) tax rate and thus reduces utility of those working, and reduces the optimal per capita-transfer and thus utility of non-workers. Since utility of all groups falls and the group whose share increases (non-workers) has lower utility than the group whose share decreases (the ℓ -workers), overall social welfare $W(m)$ monotonically decreases in m under Assumption 1.

B.2.2 Per Capita Welfare

Denote by $W^*(m) = W(\tau^*, m)$ static social welfare from consumption-labor allocation under the optimal tax-transfer policy, and define

$$W^*(m) = x\widetilde{W}^*(\mu^w(m)) \quad (51)$$

where $\widetilde{W}^*(\mu^w(m))$ is per capita welfare, which only depends on mitigation m through the share of workers in the economy. We now derive the exact form of $\widetilde{W}^*(\mu^w(m))$ and relate it to per-capita welfare under a tax-transfer-system with nondistortionary taxes. For this, note that:

$$\begin{aligned}
W^*(m) &= x^{wb}U^b(\mu^w(m)) + x^{w\ell}(m)U^\ell(\mu^w(m)) + x^n(m)U^n(\mu^w(m)) \\
&= x^w(m)U^b(\mu^w(m)) + x^{w\ell}(m)(1 + \chi)\log(1 - \xi_c) + x^n(m)U^n(\mu^w(m)) \\
&= x^w(m)U^b(\mu^w(m)) + [x^w(m) - x^{wb}](1 + \chi)\log(1 - \xi_c) + x^n(m)U^n(\mu^w(m)) \\
&= x[\mu^w(m)U^b(\mu^w(m)) + [\mu^w(m) - \mu^{wb}](1 + \chi)\log(1 - \xi_c) + (1 - \mu^w(m))U^n(\mu^w(m))] \\
&= x\widetilde{W}^*(\mu^w(m))
\end{aligned}$$

with $U^n(\mu^w(m))$, $U^b(\mu^w(m))$ given in equations (48) and (49). Exploiting these expressions we can write

$$\begin{aligned}
\widetilde{W}^*(\mu^w(m)) &= [\mu^w(m) - \mu^{wb}] (1 + \chi) \log(1 - \xi_c) + \mu^w(m) U^b(\mu^w(m)) + (1 - \mu^w(m)) U^n(\mu^w(m)) \\
&= [\mu^w(m) - \mu^{wb}] (1 + \chi) \log(1 - \xi_c) + \mu^w(m) [U^b(\mu^w(m)) - U^n(\mu^w(m))] + U^n(\mu^w(m)) \\
&= [\mu^w(m) - \mu^{wb}] (1 + \chi) \log(1 - \xi_c) \\
&\quad + \mu^w(m) \left[\log \left(\frac{\chi + \mu^w(m)}{1 + \chi} \right) - \log \left[\left(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m) \right) \right] \right] \\
&\quad - \mu^w(m) (1 + \chi) \log(1 - \xi_c) \\
&\quad + \chi \log(\chi + \mu^w(m)) + (1 + \chi) \log \left(\frac{1 - g}{1 + \chi} \right) + (1 + \chi) \log(1 - \xi_c) \\
&\quad + \log \left[\left(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m) \right) \right] \\
&= \widetilde{W}^1(\mu^w(m)) + \widetilde{W}^2(\mu^w(m))
\end{aligned}$$

where

$$\begin{aligned}
\widetilde{W}^1(\mu^w(m)) &= -\mu^{wb} (1 + \chi) \log(1 - \xi_c) + \chi \log(\chi + \mu^w(m)) \\
&\quad - \mu^w(m) \log \left(\frac{1 + \chi}{\chi + \mu^w(m)} \left(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m) \right) \right) \\
\widetilde{W}^2(\mu^w(m)) &= (1 + \chi) \log \left(\frac{1 - g}{1 + \chi} \right) + (1 + \chi) \log(1 - \xi_c) \\
&\quad + \log \left[\left(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m) \right) \right]
\end{aligned}$$

This decomposition is useful since we will show next that in the economy with lump-sum transfers and taxes per capita welfare will satisfy (see equation 54 below)

$$\frac{\partial \widetilde{W}^{LS}(\mu^w(m))}{\partial \mu^w(m)} = \frac{\partial \widetilde{W}^2(\mu^w(m))}{\partial \mu^w(m)} = \frac{1}{\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m)} > 0. \quad (52)$$

Finally we note that

$$\begin{aligned}
\frac{\partial \widetilde{W}^1(m)}{\partial \mu^w(m)} &= \frac{\chi}{\chi + \mu^w(m)} - \log \left(\frac{1 + \chi}{\chi + \mu^w(m)} \left(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m) \right) \right) \\
&\quad - \mu^w(m) \frac{\frac{(1+\chi)(\chi + \mu^w(m)) - (1+\chi)(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m))}{(\chi + \mu^w(m))^2}}{(1+\chi)(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m))} \\
&= \frac{\chi}{\chi + \mu^w(m)} - \log \left(\frac{1 + \chi}{\chi + \mu^w(m)} \left(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m) \right) \right) \\
&\quad - \mu^w(m) \frac{\chi - \mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1]}{(\mu^{bw} [(1 - \xi_c)^{-(1+\chi)} - 1] + \mu^w(m)) (\chi + \mu^w(m))}
\end{aligned}$$

To sign this derivative, we invoke Assumption 1 which guarantees

$$\frac{1+\chi}{\mu^w + \chi} \left(\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w \right) < 1$$

Then

$$\begin{aligned} \frac{\partial \widetilde{W}^1(m)}{\partial \mu^w(m)} &> \frac{\chi}{\chi + \mu^w(m)} + \mu^w(m) \frac{\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] - \chi}{(\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m)) (\chi + \mu^w(m))} \\ &= \frac{\chi}{\chi + \mu^w(m)} \left(1 + \mu^w(m) \frac{\left[\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] - \chi \right]}{\chi (\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m))} \right) \\ &\geq \frac{\chi}{\chi + \mu^w(m)} \left(1 + \mu^w(m) \frac{\left[(\mu^{bw} = 0) \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] - \chi \right]}{\chi (\mu^{bw} = 0) \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m)} \right) \\ &\geq \frac{\chi}{\chi + \mu^w(m)} \left(1 + \mu^w(m) \frac{[0 - \chi]}{\chi (0 + \mu^w(m))} \right) = \frac{\chi}{\chi + \mu^w(m)} \left(1 - \frac{\mu^w(m)\chi}{\chi \mu^w(m)} \right) = 0 \end{aligned}$$

This concludes the proof of the second part of proposition stating that

$$\frac{\partial \widetilde{W}^*(\mu^w(m))}{\partial \mu^w(m)} = \frac{\partial \widetilde{W}^1(\mu^w(m))}{\partial \mu^w(m)} + \frac{\partial \widetilde{W}^2(\mu^w(m))}{\partial \mu^w(m)} > \frac{\partial \widetilde{W}^{LS}(\mu^w(m))}{\partial \mu^w(m)} > 0. \quad (53)$$

B.2.3 Economy with Lump-Sum Transfers

Consider the problem of a government that can levy lump-sum taxes to redistribute resources. Households choose labor supply optimally, and given that there are no distortions, hours worked equal to the efficient level

$$\begin{aligned} h^b &= 1 \\ h^\ell &= (1 - \xi_c)^\chi \end{aligned}$$

and from the household budget constraints

$$\begin{aligned} c^b &= 1 + T^b \\ c^\ell &= (1 - \xi_c)^{1+\chi} + T^\ell \\ c^n &= T^n \end{aligned}$$

with associated utilities

$$\begin{aligned} U^b &= \log \left(\frac{1}{1+\chi} + T^b \right) \\ U^\ell &= \log \left(\frac{(1 - \xi_c)^{1+\chi}}{1+\chi} + T^\ell \right) \\ U^n &= \log (T^n) \end{aligned}$$

The government objective is

$$\begin{aligned} W^{LS}(m) &= x^{bw} U^b + x^{\ell w}(m) U^\ell + x^n(m) U^n \quad s.t. \\ 0 &= x^{bw} T^w + x^{\ell w}(m) T^\ell + x^n(m) U^n + gY \end{aligned}$$

with first order conditions

$$\begin{aligned} \frac{x^{bw}}{\frac{1}{1+\chi} + T^b} &= \lambda x^{bw} \\ \frac{x^{\ell w}(m)}{\frac{(1-\xi_c)^{1+\chi}}{1+\chi} + T^\ell} &= \lambda x^{\ell w}(m) \\ \frac{x^n(m)}{T^n} &= \lambda x^n(m) \end{aligned}$$

and thus

$$\lambda = \frac{1}{\frac{1}{1+\chi} + T^b} = \frac{1}{\frac{(1-\xi_c)^{1+\chi}}{1+\chi} + T^\ell} = \frac{1}{T^n}$$

and therefore

$$\begin{aligned} T^n &= \frac{1}{1+\chi} + T^b = \frac{(1-\xi_c)^{1+\chi}}{1+\chi} + T^\ell \\ 0 &= x^{bw} T^w + x^{\ell w}(m) T^\ell + x^n(m) T^n + gY^{LS}(m) \\ W^{LS}(m) &= x \log(T^n) \end{aligned}$$

where

$$Y^{LS}(m) = x^{bw} + x^{\ell w}(m)(1-\xi_c)^{1+\chi}$$

is undistorted real output (simply a function of undistorted labor supply in both sectors), as emerging in the economy with lump-sum transfers. Now note

$$\begin{aligned} T^b &= T^n - \frac{1}{1+\chi} \\ T^\ell &= T^n - \frac{(1-\xi_c)^{1+\chi}}{1+\chi} \end{aligned}$$

and the government budget constraint becomes

$$\begin{aligned} 0 &= x^{bw} \left(T^n - \frac{1}{1+\chi} \right) + x^{\ell w}(m) \left(T^n - \frac{(1-\xi_c)^{1+\chi}}{1+\chi} \right) + x^n(m) T^n + gY^{LS}(m) \\ &= xT^n + gx^{bw} + gx^{\ell w}(m)(1-\xi_c)^{1+\chi} - \frac{x^{bw}}{1+\chi} - \frac{x^{\ell w}(m)(1-\xi_c)^{1+\chi}}{1+\chi} \\ xT^n &= \left(\frac{1}{1+\chi} - g \right) [x^{bw} + x^{\ell w}(m)(1-\xi_c)^{1+\chi}] \\ T^n &= \left(\frac{1}{1+\chi} - g \right) \frac{x^{bw} + x^{\ell w}(m)(1-\xi_c)^{1+\chi}}{x} \end{aligned}$$

$$= \left(\frac{1}{1+\chi} - g \right) \frac{Y^{LS}(m)}{x}$$

Furthermore, since utility of all households is equalized, static social welfare is given by

$$W^{LS}(m) = x \log(T^n) = x \log \left[\frac{1}{1+\chi} - g \right] + x \log \left(\frac{Y^{LS}(m)}{x} \right) = x \widetilde{W}^{LS}(m)$$

and equal to a piece that captures the distortions on labor supply $\frac{1}{1+\chi}$ and the loss from government purchases. The second piece is simply equal to the number of people x time utility from per-capita income $\frac{Y^{LS}(m)}{x}$.

Now note that

$$\begin{aligned} \frac{Y^{LS}(m)}{x} &= \frac{x^{bw} + x^{\ell w}(m)(1 - \xi_c)^{1+\chi}}{x} \\ &= \frac{x^{bw} [1 - (1 - \xi_c)^{1+\chi}] + [x^{\ell w}(m) + x^{bw}] (1 - \xi_c)^{1+\chi}}{x} \\ &= (1 - \xi_c)^{1+\chi} \left[\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m) \right] \end{aligned}$$

and thus

$$\begin{aligned} \widetilde{W}^{LS}(m) &= \log \left[\frac{1}{1+\chi} - g \right] + (1 + \chi) \log(1 - \xi_c) \\ &\quad + \log \left(\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m) \right) \end{aligned}$$

and thus depends on m only through the share of working individuals $\mu^w(m)$. Furthermore

$$\frac{\partial \widetilde{W}^{LS}(m)}{\partial \mu^w(m)} = \frac{1}{\mu^{bw} \left[(1 - \xi_c)^{-(1+\chi)} - 1 \right] + \mu^w(m)} > 0 \quad (54)$$

C Flow Utilities for Each Household Type

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

$$\begin{aligned} W^\ell(x, m) &= \frac{(x^{y\ell n} + x^{y\ell e} + x^{y\ell r})}{x^\ell} [(1 - m)U^\ell(m) + mU^n(m) + \bar{u}] \\ &\quad + \frac{(x^{y\ell f} + x^{y\ell e})}{x^\ell} [U^n(m) + \bar{u} - \hat{u}] \\ W^b(x, m) &= \frac{(x^{ybn} + x^{ybe} + x^{ybr})}{x^b} [U^b(m) + \bar{u}] + \frac{(x^{ybf} + x^{ybe})}{x^b} [U^n(m) + \bar{u} - \hat{u}] \\ W^o(x, m) &= U^n(m) + \bar{u} - \frac{(x^{yof} + x^{yoe})}{x^o} \hat{u}. \end{aligned}$$

D Estimating Elasticity of Utility Weights

In order to estimate the effect of a rise in hospitalizations on the relative weight of luxury goods, we start by creating price indices for luxury and basic goods in the Consumer Price Index. Our basic goods comprise food at home, fuel oil, electricity, utilities, used vehicles, shelter, and medical services while luxury goods comprise food away from home, gasoline, new vehicles, apparel, and all remaining services and commodities not categorized as basic. We use these price indices to create the relative price of luxury to basic goods from January through December 2020 and normalize December to one.

We then estimate the following regression implied by the theoretical model:

$$\log \frac{P_t^{lux}}{P_t^{basic}} = \eta Hospitalized_t + \varepsilon_t, \quad (55)$$

where $Hospitalized_t$ is the average fraction of people hospitalized for COVID-19 in month t .

Our estimated effect of hospitalizations on relative prices, and therefore the relative utility weight on luxury goods, is $\hat{\eta} = 156.5$ with a Newey-West standard error of 48.9. The predicted time series for relative prices correlates well with the data, as illustrated in Figure 12.

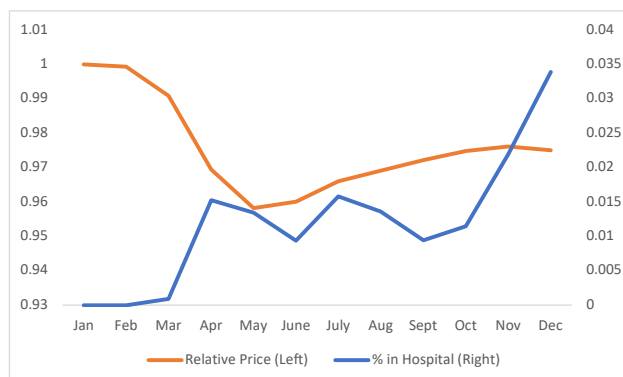


Figure 12: Relative Price and Hospitalizations

Table 8: Millions of People in Each Health State

	<i>S</i>	<i>A</i>	<i>F</i>	<i>E</i>	<i>R</i>	<i>D</i> (1000's)
03/21/20	326.37	1.99	0.71	0.02	0.91	1.32
04/12/20	320.31	1.35	1.33	0.08	6.91	27.00
06/29/20	303.84	0.91	0.89	0.06	24.17	132.85
09/30/20	289.18	0.93	0.79	0.04	38.86	207.50
12/31/20	243.42	3.81	3.48	0.21	78.72	364.75
12/31/21	191.34	0.00	0.00	0.00	138.00	657.47

E Additional Results and Figures

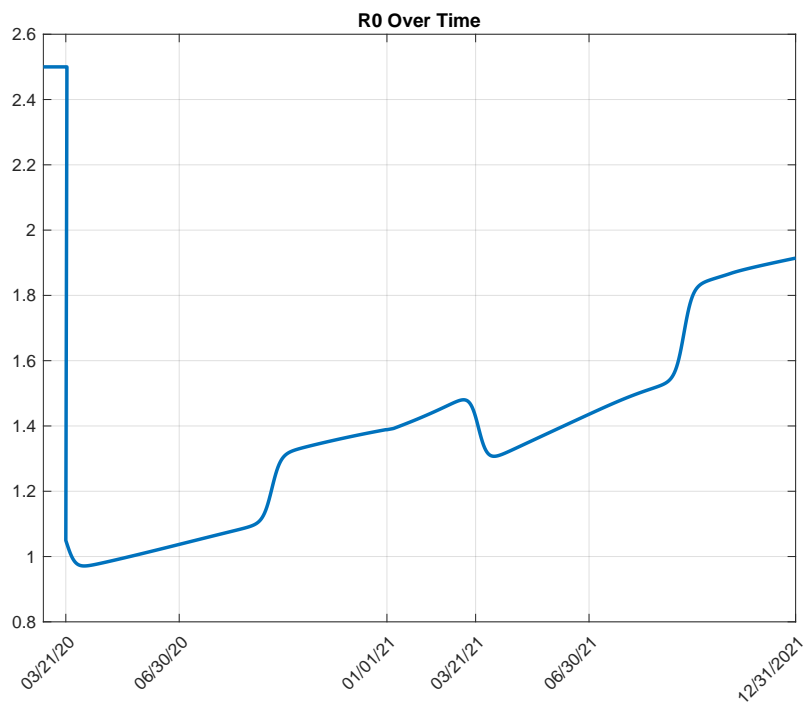


Figure 13: Model-Implied R_0

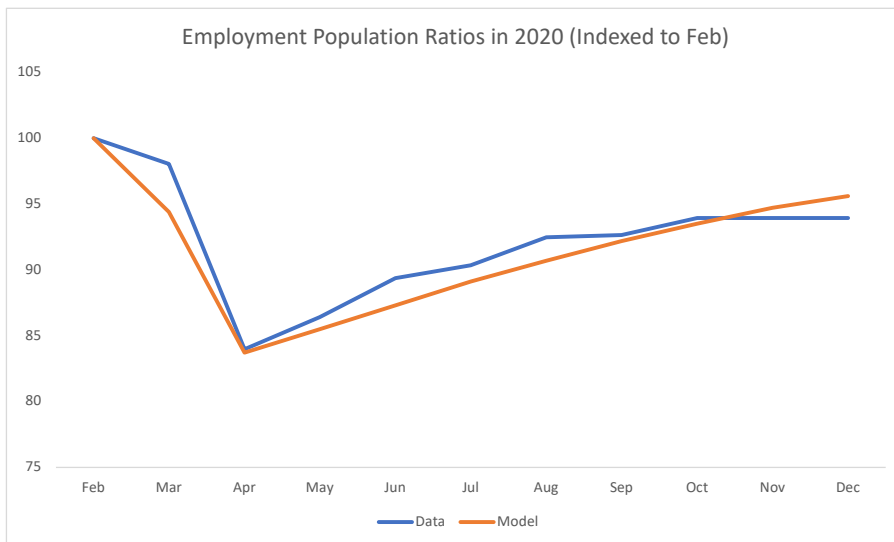


Figure 14: Employment-Population Ratio: Predicted vs Data

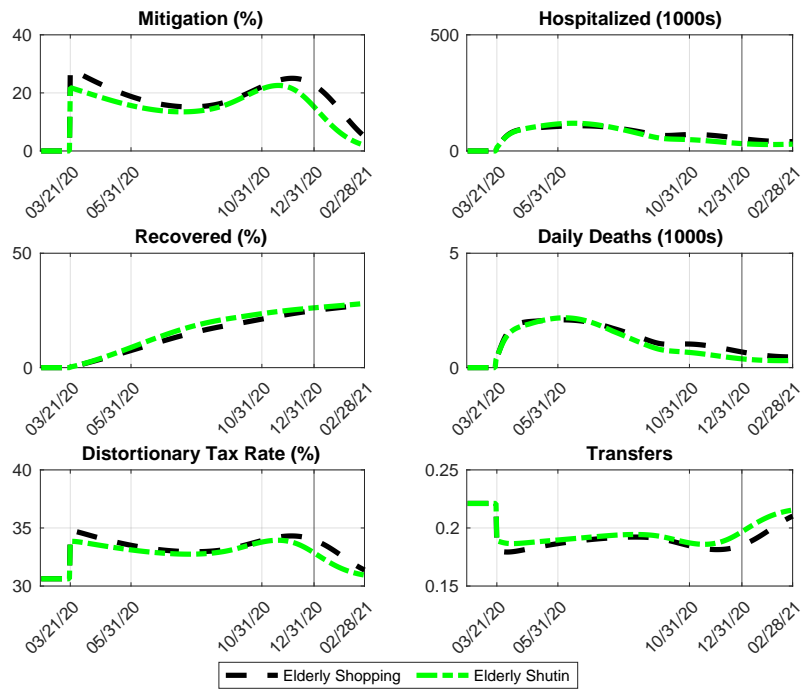


Figure 15: Key Outcomes, Baseline vs. Elderly Shut In