FEDERAL RESERVE BANK OF MINNEAPOLIS
Research Division

STAFF REPORT
No. 636

Optimal Age-Based Vaccination and Economic Mitigation Policies for the Second Phase of the Covid-19 Pandemic

January 2022

Andrew Glover
Federal Reserve Bank of Kansas City

Jonathan Heathcote
Federal Reserve Bank of Minneapolis and CEPR

Dirk Krueger
University of Pennsylvania, NBER, and CEPR

DOI: https://doi.org/10.21034/sr.636
Keywords: COVID-19; Vaccination paths
JEL classification: E21, E63

Optimal Age-Based Vaccination and Economic Mitigation Policies for the Second Phase of the Covid-19 Pandemic*

Andrew Glover† Jonathan Heathcote‡ Dirk Krueger§

January 15, 2022

First version, September 7, 2021

Abstract

In this paper, we ask how to best allocate a given time-varying supply of vaccines across individuals of different ages during the second phase of the Covid-19 pandemic . Building on our previous heterogeneous household model of optimal economic mitigation and redistribution (Glover et al., 2021), we contrast the actual vaccine deployment path, which prioritized older, retired individuals, with one that first vaccinates younger workers. Vaccinating the old first saves more lives but slows the economic recovery, relative to inoculating the young first. Vaccines deliver large welfare benefits in both scenarios (relative to a world without vaccines), but the old-first policy is optimal under a utilitarian social welfare function. The welfare gains from having vaccinated the old first are especially significant once the economy is hit by a more infectious Delta variant in the summer of 2021.

Keywords: COVID-19; Vaccination Paths

JEL Codes: E21, E63

*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. We thank our discussants Cezar Santos and Luiz Brotherhood for many useful and constructive comments.

†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

1 Introduction

What is the best way to allocate a given time-varying supply of vaccines across individuals of different ages during the second phase of the Covid-19 pandemic, starting from January 2021? To answer this applied policy question, we use the model constructed in Glover et al. (2021) which integrates an epidemiological SIR type-model into an economic model with household heterogeneity between young essential workers, young nonessential workers and old individuals.

Our previous paper did not analyze age-based vaccination policies. However, the health benefits and economic costs of vaccinations and non-pharmacological interventions such as economic shutdowns accrue to very different groups of the population, even when the possibility of transfers (financed by distortionary taxes) is considered. It is therefore natural to consider age-based health and economic mitigation policies. The primary example of an age-based policy is the prioritization of vaccines for older individuals.

In a first step, we simulate the evolution of the Covid-19 health and economic crisis from early 2021 on, based on the current vaccination plan and under alternative assumptions for economic mitigation policies. The goal is to accurately predict the second wave of infections and deaths in early 2021 and to understand the extent to which the emergence of a third wave in the spring and summer of 2021 depends on the relaxation of economic mitigation measures, given the currently planned path of vaccinations in the U.S.

In a second step, we characterize the optimal vaccine roll-out across the three groups of the population, given an exogenously fixed time path for the overall supply of vaccines and an exogenous path of economic mitigation. Current policy prioritizes vaccination of the old, since these individuals are most vulnerable to severe illness and death from Covid-19. On the other hand, since the old are not typically economically active, vaccinating younger individuals first reduces the loss of output from infections in the workplace, increasing output and thus the capacity of the economy to pay for transfers to those whose jobs are shuttered.

In a third step, we study the optimal combination of economic lockdowns and vaccinations.

Our objective is to evaluate whether opening up the economy based on the announcement of massive vaccination programs is optimal, or whether a complete end of the lockdown should wait until a majority of the working-age U.S. population is vaccinated. Furthermore, since the increase in infections associated with opening up the economy depends on whether workers or retirees are vaccinated first, we will study whether the conclusions from the second step are reversed once jointly optimal vaccination and economic lockdown policies are considered.

We conduct our analysis within the macro-epidemiological model of a health pandemic with heterogeneous households constructed by Glover et al. (2021), who study the interaction between macro-mitigation and micro-redistribution policies during the Covid-19 crisis. In this model, individuals differ by age, the sector of the economy they work in, and their health status. The population health distribution evolves according to an augmented version of the standard SIR model, with infections occurring at work, while consuming, at home through social interactions, and in the hospital when health professionals treat Covid-19 patients. The government has at its disposal three broad policy tools: (1) it can shut down part of the economy to stem the spread of Covid-19 in the workplace and during consumption activities, (2) it can provide transfers to those not working, financed by distortionary taxation, and (3) it can decide how to allocate an exogenously given (but time-varying) allotment of vaccines across different groups of the population. The benevolent Ramsey government chooses the entire time path of economic mitigation and transfers and taxes, as well as the allocation of vaccines across the young and the old, in order to maximize utilitarian social welfare, taking as given equilibrium consumption and labor supply choices by private citizens.

Our main thought experiment is threefold. First, we characterize the optimal combination of mitigation and redistribution for 2021, under the assumption that vaccines are distributed across different ages according to the actually observed vaccination rates among the young and the old in the U.S. Second, we contrast the health, economic and welfare outcomes of this baseline scenario with a scenario in which all vaccine supply is directed to the old until the old are fully vaccinated and, alternatively, one in which the young are initially prioritized. Here we hold non-pharmacological interventions (i.e., lockdowns) constant at the baseline path. Finally, we study the interaction between pharmacological and non-pharmacological interventions by letting the government optimize over the mitigation paths.

In our previous paper, we showed that relative to those of a scenario with no (hope of a) vaccine, optimal shutdown policies in the anticipation and then deployment of a large-scale vaccination campaign are more drastic in the very short run but feature more rapid reopening as the vaccine diffuses through the population. That is, vaccines and non-pharmacological interventions such as lockdowns are policy complements in the short run.

Our principal findings from the threefold new thought experiments in this paper can be characterized as follows. First, vaccinating the old first is preferred to vaccinating the young first, according to a utilitarian social welfare function. Second, this result masks substantial heterogeneity of policy preferences within the population, as all young workers (who make up 85% of the population) lose from such prioritization (relative to a young-first policy and also to our calibrated benchmark which vaccinates some young individuals early). Third, the endogenous adjustment of non-pharmacological interventions such as economic lockdowns moderates this divergence in policy preferences. If all the old are vaccinated first, the Ramsey government optimally responds by loosening the extent of lockdowns. Conversely, if the young receive the vaccine first, optimal economic mitigation is tightened in order to mitigate the associated increase in the death toll from Covid-19.

We then extend our analysis by considering the emergence of the more severe Delta variant of the virus in the summer of 2021. We show that in hindsight, an old-first vaccination policy was especially beneficial, in terms of both lives saved and economic activity, when Delta hit unexpectedly.

The paper unfolds as follows. In the next subsection, we briefly discuss the extant literature. A summary of the model is contained in Section 2; Section 3 discusses the calibration. The results are contained in Section 4, and Section 5 discusses an extension to the emergence of the Delta variant. Section 6 concludes. The appendix includes a more detailed description of the epidemiological model as well as its parameterization.

1.1 Related Literature

Our paper is related to the by now vast literature on the short-run economic impact of the Covid-19 health crisis and its optimal control through government non-pharmacological interventions. It builds on our previous work, Glover et al. (2021), as well as the closely related literature using dynamic epidemiological-economic models of the pandemic to study economic policy, mostly through the lens of (constrained) social planner problems. Representative works from this literature include Eichenbaum et al. (2021), Argente et al. (2021), Kaplan et al. (2020), Krueger et al. (2020), Guerrieri et al. (2020), Berger et al. (2022), Chari et al. (2021), Hall et al. (2020), Acemoglu et al. (2021), Boppart et al. (2020) and Brotherhood et al. (2020).

There are still relatively few papers that explicitly model the interaction between the deployment of vaccines and (optimal) economic policy. Examples include Gonzalez-Eiras and Niepelt (2020), Bognanni et al. (2020), Abel and Panageas (2020), Garriga et al. (2020), Gollier (2021) and Makris and Toxvaerd (2021).

2 The Model

The model employed in this paper is described in detail in Glover et al. (2021). It integrates an extended version of the standard epidemiological SIR model into a dynamic macro model with household heterogeneity. To keep the model focused on the dimensions of heterogeneity we think of as crucial for the Covid-19 crisis, we restrict attention to three dimensions. First, households can be young, y, or old, o. The key distinction between the young and the old is that the former group works, whereas the latter group is retired and does not participate in the labor market. In addition, the health parameters of the model differ by age, reflecting differential Covid-19 infection and death risk. The initial share of the population that is young is denoted by the parameter μ^y, and the initial share of the population that is old is correspondingly given by $\mu^o = 1 - \mu^y$.

The second source of heterogeneity is the sector (or occupation) in which a young individual works: the basic, essential sector b, in which, by assumption, production is so essential that it cannot be restricted in order to stem the pandemic, or the nonessential luxury sector, denoted by ℓ, which is subject to government lockdowns. The fraction of young workers who are attached to the basic sector is denoted by the parameter μ^b; consequently, $\mu^\ell = 1 - \mu^b$. We assume that workers cannot change sectors during the short time horizon studied in this paper; thus, the sector of work is a fixed source of heterogeneity among young individuals.

Finally, individuals of a given age and working in a specific sector can be in one of six health states, $i \in \{s, a, f, e, r, d\}$, where s stands for susceptible (to infection with the Covid-19 disease), a for infected with Covid-19 but asymptomatic, f for miserable with a fever, e for requiring emergency care, r for recovered from the disease (and therefore immune to future infections), and finally d for being dead. In the standard SIR model, all infected individuals are lumped together in an I-group; here, the a, f, and e groups all carry the virus and can infect other individuals. The distinction between the a and the f group is that the asymptomatic have no symptoms and therefore continue to go to work and shop. In doing so, they unknowingly spread the virus. Those in the f and e states neither work nor shop, but only those in the latter group occupy scarce capacity in the emergency room. Individuals who eventually die from Covid-19 progress from susceptible (s) to asymptomatic (a) to fever (f) to emergency care (e) to dead (d), but recovery (r) is possible at every stage in this chain.

Time t is continuous, and x_t denotes the size of the population at time t, using superscripts to denote measures of specific subset of the population. For example, $x_t^{y\ell a}$ is the mass of young individuals working in the nonessential luxury sector who are infected with Covid-19 but do not experience symptoms of the disease, and $x_t^{y\ell} = \sum_{i \in \{s,a,f,e,r\}} x^{y\ell i}$ is the mass of young ℓ-sector workers in period t, and so forth. Similarly, we let $\mu_t^{y\ell a}, \mu_t^{y\ell}$ and so forth denote population shares. We often suppress the dependence of population measures on time t when there is no scope for confusion, but it is understood that these population measures are functions of time.

At time $t = 0$ the total mass of the population is $x_0 = 1$. Thus, from the assumptions above it follows that at the beginning of the pandemic, the masses of the three types of individuals are given by $x_0^{yb} = \mu^y \mu^b$, $x^{y\ell} = \mu^y(1-\mu^b)$, and $x^o = (1-\mu^y)$.

2.1 The Epidemiological Block: The SAFER Model

In this section, we describe the health transitions between states $i \in \{s, a, f, e, r, d\}$, which in turn determine the dynamics of the population's health distribution. We focus the discussion on two aspects. First, the transition from the susceptible state s to the infected but asymptomatic state a, since this transitions can be affected by non-pharmacological interventions such as locking down part of the nonessential part of the economy. Let m_t denote the share of the nonessential sector that is mitigated (locked down) at time t. Second, we discuss the health state transitions induced by vaccines.

Turning to the first point, and following Glover et al. (2021), there are four sources of possible virus contagion: infections at work, infections while consuming (or shopping), infections from interactions with family and friends outside work, and, for workers in the basic sector, infections from taking care of infected patients in the emergency room. We index these four sources as $j \in \{w,\ c,\ h, e\}$, respectively. For each of these sources, the flow of new infections for a given mass of susceptible types $(x^{ybs}, x^{y\ell s}, x^{os})$ is given by the product of the number of infected people this group meets in activity j (denoted by $x_j(m_t)$), times the infection-generating rate $\beta_j(m_t)$ from a typical meeting. Both components potentially depend on the extent of economic mitigation.

The outflow from the susceptible to the asymptomatically infected state for the three groups of individuals is then given by equations (1)-(3):

$$\dot{x}^{ybs} = -\left[\beta_c(m_t)x_c + \beta_h x_h\right]\ x^{ybs} - \beta_w(m_t)x_w(m_t)\ x^{ybs} - \beta_e x_e\ x^{ybs}, \tag{1}$$

$$\dot{x}^{y\ell s} = -\left[\beta_c(m_t)x_c + \beta_h x_h\right]\ x^{y\ell s} - \beta_w(m_t)x_w(m_t)(1-m_t)x^{y\ell s} \tag{2}$$

$$\dot{x}^{os} = -\left[\beta_c(m_t)x_c + \beta_h x_h\right]\ x^{os}. \tag{3}$$

Note the asymmetry across the three groups, in that the old do not work (and therefore get infected only at home and while shopping), and while all workers risk infection on the job unless shuttered out, only basic-sector workers staff the emergency room and face infection risk there.

In the expression above, the measures of infectious individuals who potentially cause contagion with Covid-19 in the four activities are given by

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

$$x_c = x^a \tag{5}$$

$$x_h = x^a + x^f \tag{6}$$

$$x_e = x^e. \tag{7}$$

These measures are based on the assumptions that symptomatic individuals do not work or shop by themselves and that basic and luxury sector workers can meet in the workplace. Whereas the infection rates at home and in the emergency room (β_h, β_e) are exogenous parameters, those associated with working and shopping are endogenous to mitigation policies and given by

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

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

where $\alpha_w = \beta_w(0)$ and $\alpha_c = \beta_c(0) =$ are parameters determining infection rates in the absence of mitigation and when the entire young population in the working in the luxury sector is still healthy (i.e., when $x^{\ell w} = (1 - \mu^b)\mu^y$). In Glover et al. (2021) we provide a micro foundation of these lockdown-contingent infection rates. The remainder of the epidemiological block then simply describes the mechanical transitions of individuals through the health states (asymptomatic, fever-suffering, hospitalized, and recovered) once infected.[1]

Finally, transition to death occurs only from the emergency room state at age-dependent rates $\sigma^{yed} + \varphi(x^e)$ and $\sigma^{oed} + \varphi(x^e)$, where we permit excess mortality at rate φ when hospital

[1]For completeness, we include these transitions in equations (24)–(35) in Appendix A, with parameters that are allowed to vary by age.

capacity is exceeded. This additional mortality takes the functional form

$$\varphi(x^e) = \lambda_o \max\{x^e - \Theta, 0\}. \tag{10}$$

Given emergency room capacity Θ, the term $x^e - \Theta$ in the max operator gives the extent of hospital overuse, with parameter λ_o controlling by how much the death rate of the hospitalized rises once capacity is exceeded.

Finally, we assume that both susceptible and recovered individuals receive vaccines and that the former group makes an automatic, instantaneous transition to the recovered state upon vaccination.

2.2 The Economic Model

We now describe the economic block of the model.

2.2.1 Technology

The production technology is linear in labor input in both sectors. Since the basic sector is never shuttered, output in this sector is given by the measure $x^{bw} = x^{ybs} + x^{yba} + x^{ybr}$ of young workers in this sector, times the number of hours h^b they work:

$$Y^b = x^{bw} h^b = \left[x^{ybs} + x^{yba} + x^{ybr}\right] h^b. \tag{11}$$

The underlying assumption is that asymptomatic individuals continue to work, whereas those with fever symptoms stay at home. Output in the nonessential luxury sector is defined symmetrically but depends on the extent of mitigation policy:

$$Y^\ell(m_t) = (1 - m_t) x^{\ell w} h^\ell = (1 - m_t)\left[x^{y\ell s} + x^{y\ell a} + x^{y\ell r}\right] h^\ell. \tag{12}$$

We treat the good produced in the basic sector as the numeraire, and the price of the ℓ good is denoted by p. GDP in any given period is then given by

$$Y = Y^b + pY^\ell. \tag{13}$$

2.2.2 Preferences and Endowments

Households have preferences defined over consumption and hours worked $\{c_t, h_t\}$, and they also value being alive and not being sick. Lifetime utility for a young worker is given by

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

where ρ denotes the discount rate, T^y is the remaining life span in the absence of 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 utility from consumption and labor supply, $u(c_t^y, h_t^y)$, the flow value of being alive, $\overline{u}$, and a utility discount $\widehat{u}_t^i$ for being sick.[2] Dead individuals have instantaneous utility normalized to zero. Old individuals have the same preferences but 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.

As noted above, workers who have a fever or are in hospital do not work. In equilibrium, expected utility of a young individual is sector-specific because sectors differ in the share of economic activity being shut down (and thus, in the probability of being able to work when healthy), because the distribution of health outcomes is sector-specific, and because wages in the two sectors will differ, as described below.

We assume that the period utility function takes a Greenwood, Hercowitz and Huffman form:

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

where 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. In the full dynamic model ξ_c will be a function of the state of the pandemic, but for the exposition of the static household decision problem, we can treat it as a parameter.

[2] We 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 yields disutility, and being in the emergency room yields even higher disutility.

2.3 Household Maximization

In this section, we characterize the static household maximization problem over consumption and labor, as we abstract from intertemporal asset accumulation in this paper. Since one unit of labor produces one unit of output and the basic good is the numeraire, the wage rate in the basic sector is $w^b = 1$. Similarly, because of the linear technology in the luxury sector, one unit of labor produces output worth p. A perfectly competitive labor market then implies that $w^\ell = p$. In any equilibrium in which both goods are produced and consumed in positive amounts and thus both labor markets clear, the price of luxury goods is given by

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

which is the only price at which households purchase both goods in positive amounts. Finally, we assume that the government implements a simple tax-transfer system that taxes labor income at a flat rate τ and provides a transfer T to everybody not working.

The household maximization problem of those individuals who can work (i.e., the healthy or asymptomatic young people who are not mitigated) in sector $i \in \{b, \ell\}$ and given wages $w^b = 1$, $w^\ell = 1 - \xi_c$, is

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

with solution

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

For non-working households (the old, and the young who are sick or mitigated), the budget constraint and period utility are given by

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

2.4 Aggregation and Market Clearing

The government, in addition to raising taxes and paying transfers T (all denoted in terms of the basic good), purchases goods in both sectors of the economy. We assume that these purchases are a constant share g of output in both sectors; that is, $G^i = gY^i$ for $i \in \{b, \ell\}$. The market clearing conditions then read as

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

2.5 Fiscal Policy

The government chooses the path of mitigation (shutdowns) m_t and redistribution through proportional taxes τ_t, which finance lump-sum transfers to individuals who do not or cannot work T_t. We assume below that when setting taxes and transfers, the government values all individuals equally, corresponding to equal Pareto weights in the social welfare function. Under this assumption it is optimal to equalize transfers across all non-working households. The static government budget constraint reads as

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

which determines per capita transfers to non-working households as a function of the current tax rate τ and the extent of mitigation in the luxury sector:

$$T = \frac{(\tau - g)\left[Y^b + pY^\ell(m)\right]}{x^n(m)} = T(\tau, m), \tag{17}$$

where $x^n(m)$ is the measure of non-working households (which is solely a function of mitigation m and the predetermined population health distribution).

Note that we have constructed the model in such a way that the tax-transfer policy has no impact on any health transitions. Therefore, at each date t, we can solve a static optimal tax-transfer policy problem that maximizes instantaneous social welfare, taking as given the current level of mitigation $m_t = m$ and the current population health distribution.

Period utilities for non-working individuals and for those working in the basic and luxury sectors are given by

$$U^n(\tau, m) = \log(T) = \log(\tau - g) + \chi \log(1-\tau) + \log\left[\frac{Y^{LS}(m)}{x^n(m)}\right], \tag{18}$$

$$U^b(\tau) = -\log(1+\chi) + (1+\chi)\log(1-\tau), \tag{19}$$

$$U^\ell(\tau) = -\log(1+\chi) + (1+\chi)\log(1-\tau) + (1+\chi)\log(1-\xi_c) = U^b + (1+\chi)\log(1-\xi_c). \tag{20}$$

Static social welfare in turn is given by

$$W(\tau, m) = x^{bw}U^b(\tau) + (1-m)x^{\ell w}(m)U^\ell(\tau) + x^n(m)U^n(m,\tau). \tag{21}$$

In Glover et al. (2021) we show that the optimal redistribution policy is given by

$$\tau^*(m) = (1-g)\,\frac{\mu^n(m)}{1+\chi} + g, \tag{22}$$

$$T^*(m) = \left(\frac{1-g}{1+\chi}\right)^{1+\chi} [\chi + \mu^w(m)]^\chi \; \frac{Y^{LS}(m)}{x}, \tag{23}$$

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 that is working. Thus, more mitigation, by reducing the share of the population that is working and increasing the share that is not working, translates into higher optimal tax rates and lower optimal transfers.

By inserting these expressions for the optimal tax-transfer system into the period utilities, one can construct static social welfare $W^*(m)$. In our companion paper, we show that as long as the preference shock for luxury consumption is not too large,[3] static social welfare can be decomposed as $W^*(m) = x\widetilde{W}^*(\mu^w(m))$, and per capita welfare $\widetilde{W}^*(\mu^w(m))$ is strictly increasing in $\mu^w(m)$ and thus strictly decreasing in mitigation m. Furthermore, we also demonstrate that social welfare is more elastic to mitigation in this model with distortionary taxation than in an economy in which the government has access to lump-sum taxation.

The optimal mitigation path then maximizes the discounted sum of static period welfare, taking into account the (dis-)utility flows from being alive and sickness. Its key trade-off is that a marginal increase in mitigation m entails static economic costs 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.

3 Calibration

In Glover et al. (2021), we provide a detailed discussion of the calibration of the model. Here, we summarize the key choices we make; the complete list of parameters and their calibrated values is contained in Section B of the Appendix.

Household Preferences The population is composed of $\mu^y = 85\%$ young and $\mu^o = 15\%$ old individuals who discount the future at 3% annually and have a unit Frisch elasticity, $\chi = 1$. Their remaining life expectancies are $T^y = 47.8$ and $T^o = 14.0$ years, respectively. A key preference parameter is the flow value of life $\overline{u}$, which we set so that the value of a statistical life is \$6.65 million, which, with an average remaining lifetime of 42.6 years and a discount rate of 3%, amounts to an annual flow payment of \$276,700. The disutility of fever is set to $\hat{u}^f = -0.3\left(\ln(\overline{c}) + \overline{u}\right)$, as argued by Hong et al. (2018), and we set the flow value of being in hospital equal to the flow value of being dead (zero).

The preference shifter ξ_c on luxury consumption ℓ is meant to depress demand in this sector when Covid-19 infection risk is high. We approximate this relationship with the functional form

[3] This assumption guarantees that workers in the luxury sector prefer to work rather than not work and receive transfers.

$\xi_{c,t} = 1 - \exp(\eta x_t^e)$, where x_t^e is the number of people hospitalized with Covid-19 and in the emergency room. Using the model relationship that the price for nonessential luxury goods is $p_t = 1 - \xi_{c,t}$, we estimate the parameter η by regressing the log-price on the share of the Covid-19-hospitalized population. In our companion paper, we show that this simple functional form fits the negative empirical relationship well.

Technology We assume that the luxury sector accounts for $\mu^\ell = 45\%$ of the economy.[4] The government absorbs a share of 24.7%, the share of government outlays (less social security and Medicare payments) in GDP in 2019. Given these values, the pre-Covid (that is, zero mitigation) tax rate is $\tau = 30\%$.

Medical Parameters The transition rates between the different health states (apart from the transitions from infected to asymptomatically infected) are described in detail in Glover et al. (2021). They are permitted to depend on age (young vs. old) and are chosen to match six (age-specific) empirical targets: the average length of time individuals spend in the asymptomatic, fever, and emergency-care states, and the relative chance of recovery (relative to disease progression) in each of the three states. Our choices imply initial (at the beginning of the pandemic, before March 21, 2020) infection fatality rates of 0.146% for the young and 5.84% for the old.[5]

Given these parameters, the infection parameters α_w, α_c, β_h, and β_e control the speed at which infections grow over time. We set the hospital infection-generating rate β_e so that this channel of infections accounts for 5% of cumulative Covid-19 infections though April 12, 2020, based on data on infections for healthcare workers. The parameters α_w, α_c, and β_h are chosen so that the model reproduces an initial reproduction number $R_0 = 2.5$ before March 21, 2020, and the relative shares of disease transmission that occur at work, via market consumption, and in non-market interactions.[6] Mossong et al. (2008) argue that in normal (pre-pandemic)

[4]Empirically, we associate this sector with food away from home, transportation services, apparel, new vehicles and gasoline, with the idea that these are sectors in which consumption involves social interaction.

[5]Finally, we assume that treatment for hospitalized Covid-19 patients improved so that the overall IFR falls from 1% (the weighted average of the IFR of the young and the old) on March 21, 2020, to 0.57% by December 31, 2020.

[6]For R_0 estimates see (CDC, 2020, Table 1)

times, 35% of potentially infectious inter-person contact occurs at work, 19% occurs during consumption activities and the rest occurs at home and other settings.

The basic reproduction rate should not be interpreted as a constant but rather as something that changes with social behavior and environmental factors such as the season. In Glover et al. (2021) we model this in a reduced form fashion by scaling all infection-generating rates (α_w, α_c, β_h, β_e) by a factor ζ_t. Before March 21, 2020, $\zeta_t = 1$. Because of a change in household behavior, once it became clear Covid-19 had arrived in the U.S., the scaling factor is assumed to fall to $\zeta_t = \zeta_H < 1$ on March 20, 2020. In addition, owing to seasonality of infection rates due to weather, in the warmer summer months there is a further fall in the scaling factor to $\zeta_t = \zeta_L < \zeta_H$. The parameters (ζ_L, ζ_H), the date of seasonal transition, and the initial share of infected individuals in early 2020 are chosen such that the model matches cumulative official deaths at four dates: April 12, May 31, October 31, and December 31, 2020. Thus, by construction, at the end of 2020 the cumulative deaths implied by the model match those in the data.[7] Finally, we assume that as vaccinations are rolled out in 2021, the scaling parameter ζ_t again rises proportionally with the share of newly recovered individuals, eventually reverting back to pre-Covid behavior.

For hospital capacity, we set the number of hospital beds to $\Theta = 100,000$, implying excess mortality commences when more than 0.03% of the population is hospitalized. The parameter choice of λ_o implies that when there are $200,000$ hospitalizations, the mortality rate in the emergency room is 25% above the normal rate.

Mitigation Policies Even though our focus is on the time period starting in 2021, when vaccines begin to be deployed, the actual time path for actual economic mitigation is an important determinant of the population health distribution at the end of 2020. We choose this path so that our model replicates the dynamics of employment from March 21, 2020 (the time states such as California, New York and Illinois started to impose business closures) until the end of 2020.

[7] This parameterization implies a decline of R_0 from 2.5 to 1.26 around March 21, 2020.

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

In our companion paper, we characterized optimal mitigation and redistribution for the early phase of the Covid-19 pandemic. In this paper, we focus on the second phase, starting at the beginning of 2021 as a vaccine becomes available, with a fairly predictable path for its deployment. We are especially interested in documenting the positive and normative consequences of alternative prioritizations of who gets the vaccine first by age, taking as given the total production and distribution of vaccines.

In our benchmark scenario, 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 U.S. in the first half of 2021. Specifically, we assume that susceptible and recovered individuals are vaccinated over the six month period between January 1 and July 1, 2021. Per day, 0.47% of the old and 0.30% of the young are vaccinated, which approximately replicates the differential pace of vaccination rates by age reported by the CDC.[8] We assume that after July 1, 2021, vaccinations of the young and old continue at a constant pace until 60% of the young and 80% of the old have been vaccinated.

Our simulations in this section start on January 1, 2021. The initial condition is the population's health distribution for that date implied by the model simulation for 2020 in Glover et al. (2021), which we argued was a good approximation of the actual health distribution at the end of 2020. Equipped with this initial condition and the deterministic time path for the number of available vaccines, we now explore the optimal combination of economic mitigation and vaccine allocation.

4.1 Optimal Mitigation under the Benchmark Vaccination Plan

To establish a benchmark for comparison, under the empirically observed vaccination plan, which gives early vaccine doses mostly to the old but also reserves some for the young, we now document the optimal mitigation path as a function of the social welfare function. Figure 1

[8] See https://covid.cdc.gov/covid-data-tracker/vaccination-demographics-trends).

shows the preferred mitigation paths in the presence of a vaccine roll-out in the first half of 2021. The figure shows that the old prefer more mitigation than the young working in either sector, and basic workers support a quicker reduction of lockdowns as the winter wave of infections subsides. The utilitarian optimum lies in between these group-specific policy preferences. In Glover et al. (2021) we argue that the optimal *level* of mitigation is significantly higher in the presence of a vaccine than in its absence, since the (anticipated) diffusion of a vaccine implies that infections prevented by mitigation will never occur. Without a vaccine, in contrast, mitigation primarily delays infections.

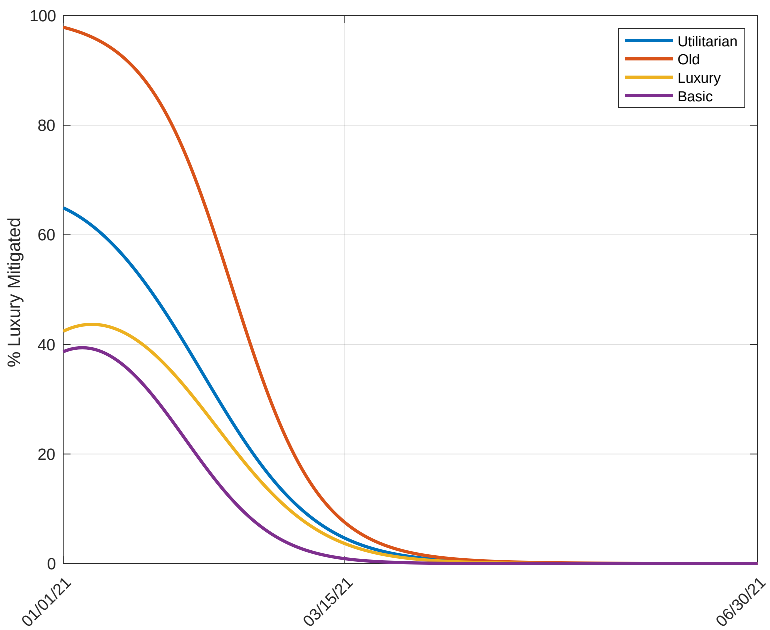

Figure 1: Preferred Mitigation in the Presence of a Vaccine

Our previous paper, Glover et al. (2021), shows that the vaccine saves 335,000 lives, with about half of these lives saved due to the vaccine itself (holding fixed the path of shutdowns). The other half is due to the fact that with a vaccine, a stronger mitigation policy response is optimal in early 2021, in turn softening the winter wave in January and February 2021. See the third row of Table 1, which we reproduce here for comparison.

Although the stronger optimal shutdowns with a vaccine imply a recession very early in 2021

Table 1: 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%

that is deeper than the one in the scenario without a vaccine, the fact that the economy can be opened up relatively quickly in early 2021 implies that cumulative output for 2021 is actually higher with a vaccine than without, under the respective optimal mitigation policies. The last row of Table 1 indicates a 0.12% boost to GDP from the combination of vaccines and the associated re-optimization of the path for mitigation.

In terms of the size of the welfare gains from obtaining a vaccine relative to those of the situation in which a vaccine never becomes available, the first row of Table 1 shows that these are large for the utilitarian Ramsey government, equivalent to increasing consumption permanently by two-thirds of 1% for individuals' remaining lifetimes. Comparing the first and the third column indicates that about half the welfare gains come from the vaccine itself, and the remaining half from the adjustment of the mitigation path to the availability of the vaccine.

The second row of the table shows that the welfare gains are very heterogeneous across the population and much larger for the old, at close to 6% of permanent remaining lifetime consumption. The main source of these gains is the sharply reduced number of deaths, especially among the elderly, again in roughly equal parts stemming from the vaccine itself and from the associated mitigation policy response, which calls for more drastic shutdowns early on and quick reopening once vaccinations have protected the old.

4.2 Alternative Vaccination Paths

We now study the health, economic and welfare implications of alternative vaccination plans, reserving initial vaccine supply either exclusively for the old or exclusively for the young. Concretely, under the old-first (young-first) scenario all vaccines are given to the old (young) until 80% (60%) of the group has been vaccinated. Allocations then switch to the young (old) until 60% of the young (80% of the old) have been vaccinated. Although the old are significantly more vulnerable to death from the disease, they are less likely to get infected and thus to spread the virus. It is therefore possible that administering the vaccine to the young first might slow the wave of infections in the massive winter wave of early 2021. As a result it is not a priori obvious that an optimally designed vaccine roll-out should have prioritized the old rather than the young.

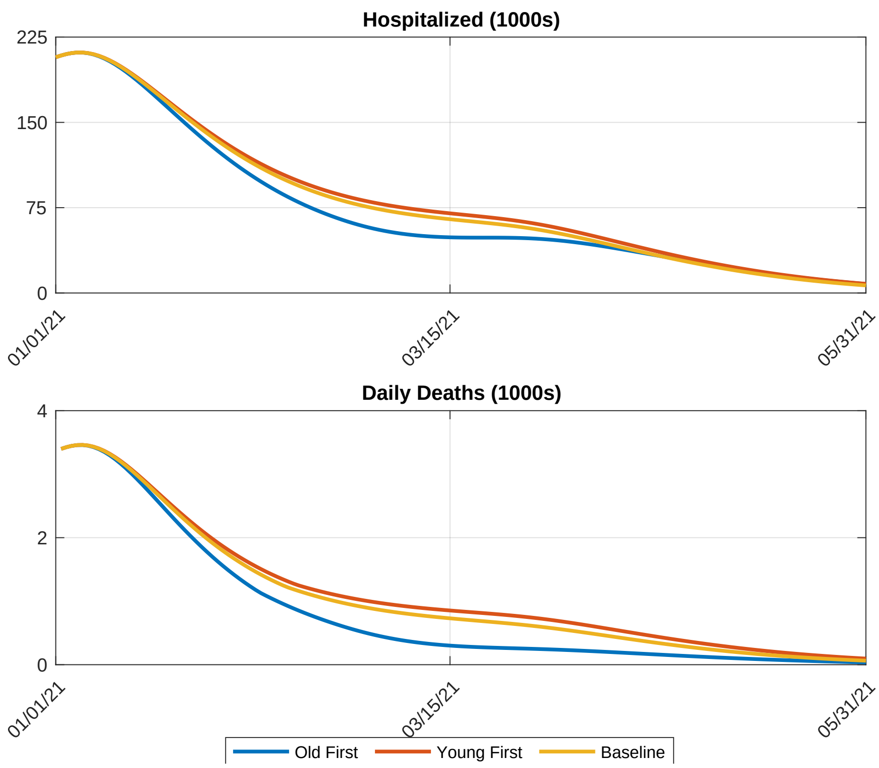

Figure 2: Vaccination of the Young versus the Old: Baseline Mitigation Policy

To conduct our analysis most transparently, in a first step, we hold economic mitigation policy constant at the old benchmark (the optimal policy associated with the baseline vaccination plan) before permitting an optimal mitigation policy response to the changing vaccine regime in the next subsection. Figure 2 displays the outcome under this first thought experiment. The figure shows that hospitalizations and deaths are substantially higher in early 2021 if the Ramsey government decides to vaccinate the young first, rather than the old.

4.3 The Interaction of Optimal Vaccination and Mitigation

In this section, we allow the Ramsey government to adjust the optimal path of economic shutdowns to the age-dependent vaccine roll-outs studied in the previous section. Figure 3 displays the *optimal* mitigation paths for the first half of 2021, under the benchmark vaccination scenario and the two alternative scenarios in which either the old or the young receive absolute priority in vaccinations.

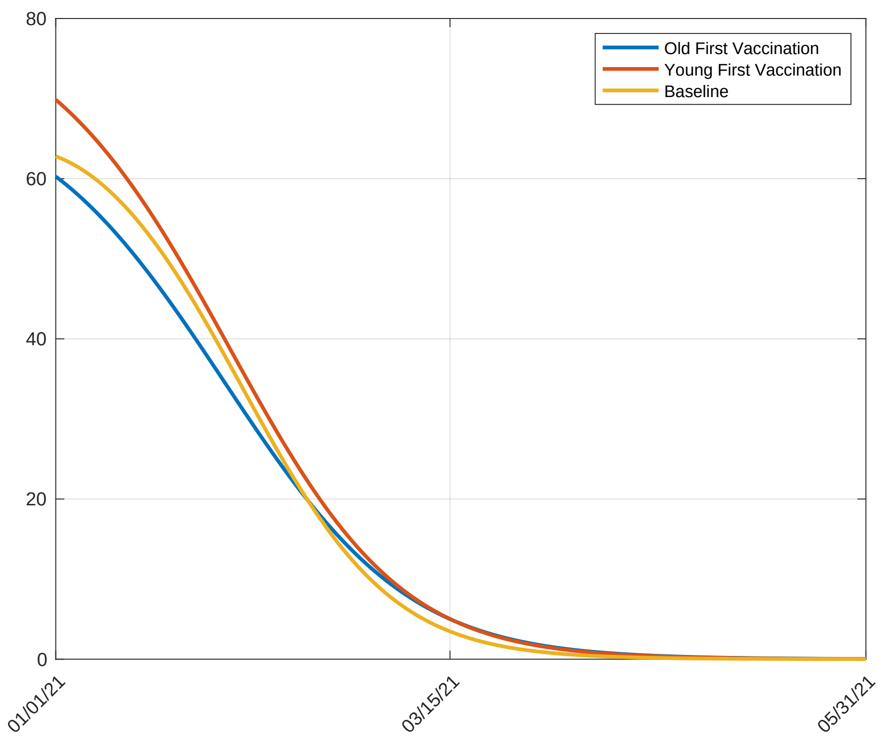

Figure 3: Optimal Mitigation under Alternative Vaccination Paths

The main observation from Figure 3 is that a utilitarian Ramsey government offsets the additional deaths threatened by a young-first policy by mitigating more early on (when the old remain unvaccinated and thus vulnerable to infection). In contrast, an old-first vaccination plan protects an increasing share of the old through vaccines, and thus shutdowns can be milder.

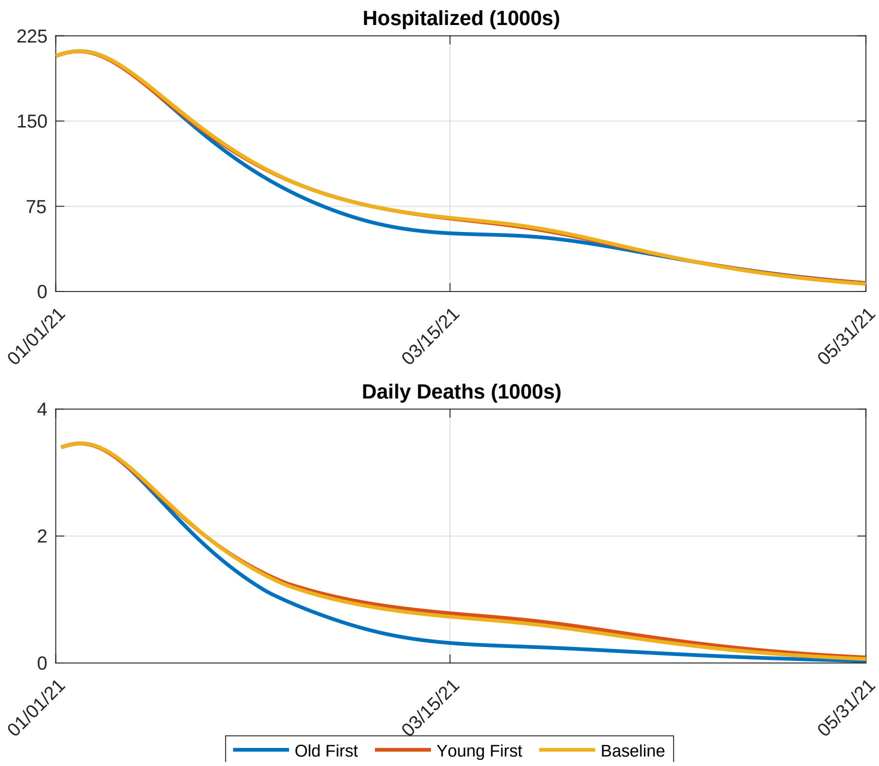

Figure 4: Vaccination of the Young versus the Old: Health Outcomes under Optimal Policy

Figure 4 shows the evolution of hospitalizations and deaths in the three different scenarios (as in Figure 2) but now permits the Ramsey government to optimize shutdowns in each vaccination scenario; that is, it applies the mitigation paths displayed in Figure 3. The stronger mitigation under the young-first policy helps close the gap in hospitalizations and deaths between the young-first- and the old-first vaccination paths but is not enough to fully offset it. As a consequence, the former vaccination strategy, even with adjustment of other policies, results in more adverse population health outcomes in the second phase of the Covid-19 pandemic, especially for the elderly population. In contrast, giving all available vaccines initially to the old would have significantly brought down hospitalizations and deaths (relative to the young-first scenario, but even relative to the benchmark path, which prioritized the old but not fully). This is especially

Table 2: Health and Economic Consequences and Welfare Gains from Vaccinating Young First

	Fixed Policy	Optimal Policy
Young Basic	0.014%	0.006%
Young Luxury	0.007%	-0.001%
Old	-0.304%	-0.206%
Utilitarian	-0.015%	-0.014%
Deaths Avoided	-12,791	-5,541
GDP Gain	-0.05%	-0.35%

Table 3: Health and Economic Consequences and Welfare Gains from Vaccinating Old First

	Fixed Policy	Optimal Policy
Young Basic	-0.045%	-0.043%
Young Luxury	-0.041%	-0.037%
Old	0.852%	0.836%
Utilitarian	0.030%	0.032%
Deaths Avoided	33,736	31,559
GDP Gain	0.06%	0.15%

true for the month of March 2021; see again Figure 4.

Tables 2 and 3 quantify these statements. They summarize the health, economic and welfare consequences from the different vaccination paths when combined with the optimal mitigation policy responses given these paths. Table 2 focuses on the young-first vaccination plan and Table 3 on the path that prioritizes the old. The point of comparison is always the benchmark vaccination path that rolls out vaccines to both age groups and the associated *optimal* mitigation path.

Focusing first on two summary measures of the health and economic consequences of the alternative vaccination strategies, we see that relative to the benchmark, a young-first policy (Table 2) leads to approximately 12,000 more deaths and roughly unchanged economic performance, holding lockdown policies constant. The Ramsey government responds to this threat of additional human carnage by tightening economic mitigation relative to the benchmark mitigation path (see again Figure 3). Thus, under the the optimal policy the additional loss of lives falls to about 5,500 extra deaths. As a consequence, relative to the benchmark, under the optimal policy response a young-first vaccination strategy leads to more deaths *and* a weaker economy, and thus lower utilitarian social welfare.

In contrast, an old-first policy (Table 3) saves extra lives, and the optimizing Ramsey government responds by loosening lockdowns more quickly for extra economic gain (and at the expense of some of the extra lives saved). As a result, under the optimal policy combination, GDP is 0.15% higher than under the benchmark, and approximately 31,500 fewer (predominately old) individuals lose their lives during the second phase of the Covid-19 pandemic.

It might at first seem surprising that vaccinating the young first does not at least benefit economic growth. Figures 5 and 6 further explore this issue by plotting, against time, (1) the relative price of output in the nonessential luxury sector, (2) the share of the young population actually working, and (3) GDP, relative to the pre-Covid steady state. The underlying economic mitigation path is held fixed across vaccination scenarios at the baseline (the optimal path under the baseline vaccination scenario) in Figure 5, while mitigation varies across scenarios in Figure 5, with mitigation paths set to the scenario-specific optima.

The first figure shows that the more dramatic winter wave of infections and hospitalizations under a young-first vaccine policy significantly reduces the price of output (and thus wages) in the luxury sector in the first half of 2021, relative to an old-first policy. This in turn depresses labor supply of workers in this sector. In contrast, although under a young-first policy slightly more individuals are working (since fewer young workers transit into the fever (f) and hospital (e) states), this effect is quantitatively rather small. It is small because most non-vaccinated young do not fall sick even if infected, and if they do, they on average spend only a few days in a state in which they cannot work. Consequently, the price effect dominates, and a young-first vaccination policy (holding mitigation policy fixed) leads to slightly lower output than an old-first policy (compare again the first columns in Tables 2 and 3), even though it shields young workers from the disease.

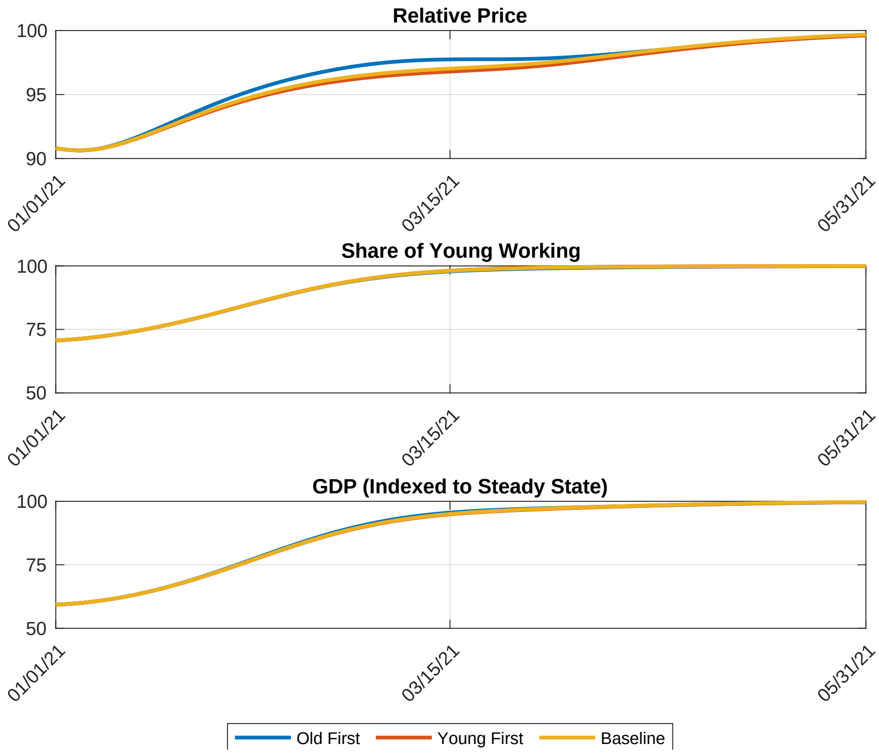

Figure 5: Economic Indicators: Holding Mitigation Fixed

This difference in output is reinforced by the optimal mitigation policy response to the different vaccination strategies. Recall that with a young-first policy, the government responds to the (threat) of more elderly deaths with more stringent economic lockdowns of the luxury sector, which depress GDP. Conversely, a reduction in deaths by first inoculating the old allows the Ramsey government to relax such measures. Figure 6 shows that the share of the young who work is significantly smaller under a young-first policy once the mitigation policy is re-optimized to that vaccination profile. The young who are not working are overwhelmingly healthy but cannot work because the government tries to protect the unvaccinated old by shutting down part of the economy. As a consequence, output is significantly lower (and deaths higher) under a young-first vaccination policy, given lockdown strategies that are optimally tailored to vaccination profiles.

In terms of welfare, Table 3 shows that the old strongly prefer to be vaccinated first. They

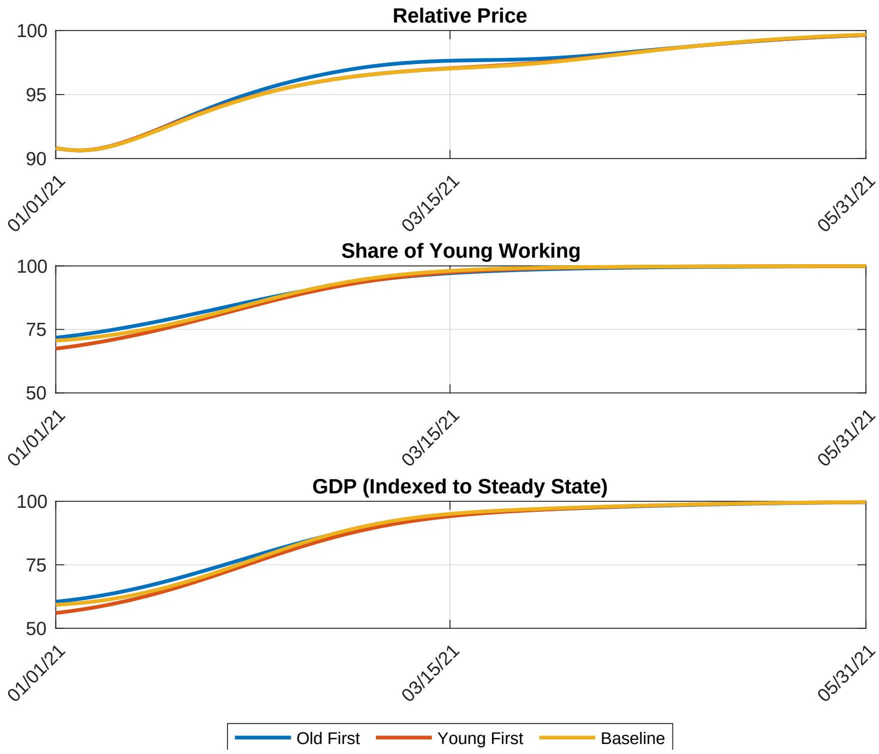

Figure 6: Economic Indicators: Optimal Mitigation

face significant welfare losses under the young-first policy (because the additional loss of lives accrues almost entirely to them), although less so if the Ramsey government responds optimally with stronger economic mitigation.

The young prefer to be vaccinated first, although their welfare consequences from different vaccination paths are quantitatively an order of magnitude smaller than those for the old. That is, the young and the old have opposite policy preferences, but there is much more at stake for the old than for the young in terms of what vaccination and economic mitigation paths the Ramsey government decides to follow.

As Tables 2 and 3 show, with the young-first policy the large group of the young see very moderate welfare gains, and the smaller group of the old see significant welfare losses. Consequently, utilitarian welfare is 0.014% lower than under the benchmark vaccination policy, which mostly vaccinates the old first.[9] In contrast, utilitarian social welfare increases under an old-first policy because of sizable welfare gains for the old and despite moderate welfare losses for young workers.

Thus, to summarize, and taking utilitarian social welfare as a normative summary measure, we conclude that a vaccination plan that exclusively prioritizes the old is optimal. It is important to note, though, that this plan is by no means a Pareto improvement over the benchmark; in fact, 85% of the population loses under this plan, although only moderately so.

5 The Emergence of the Delta Variant

We conclude this paper by investigating how different vaccination strategies fare as the economy is unexpectedly hit by a new, more infectious variant of the virus. The Delta variant of Covid-19 emerged in the U.S. around the beginning of May 2021, accounted for approximately 70% of all new infections by the beginning of July, and is responsible for virtually all new cases of Covid-19 in the U.S. at the time of writing (November 2021).[10]

To model the emergence of Delta, we take the following simple approach. We assume that on June 25, 2021, infection rates rise by a factor of three. In addition, all vaccinations throughout 2021 are only partially effective, in that only 50% of susceptible individuals who are vaccinated transition to the recovered state. The first assumption reflects the higher contagiousness of the Delta variant. The second is motivated by the observation of a significant number of infections among previously vaccinated individuals. These assumptions, in combination with our estimates of the empirical vaccination and mitigation paths, imply a third wave of deaths in the fall of 2021 in the U.S. that accords well with the data, in terms of both magnitude as well as the timing of the peak. In what follows, the economic mitigation path is set to our estimate of the empirical path, which implies very modest shutdowns in the summer and fall of 2021.

The health dynamics under different vaccination policies (and under the empirically observed

[9] The reason young luxury workers care more about health outcomes than young basic workers is that a stronger wave of hospitalizations depresses demand for their services and thus prices and wages in the luxury sector.

[10] See https://covid.cdc.gov/covid-data-tracker/.

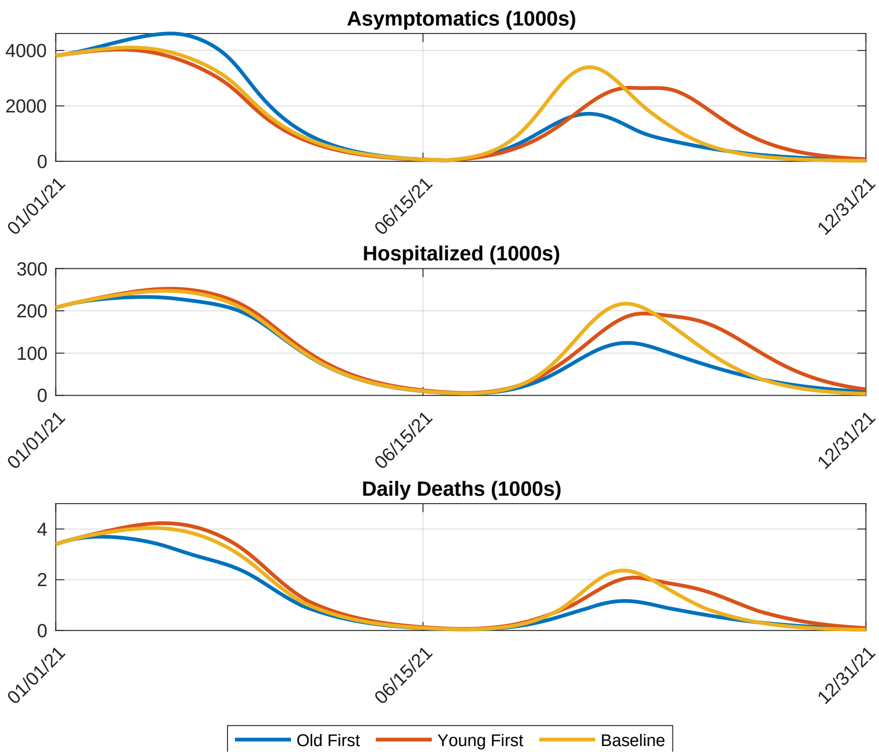

Figure 7: Vaccination of the Young versus the Old: Health Outcomes under Optimal Policy

mitigation path, which is maintained throughout this section) are displayed in Figure 7. The main difference across vaccination scenarios is how many young and old remain unvaccinated in the early summer of 2021, when the Delta variant unexpectedly arrives, and who gets prioritized for new vaccines thereafter.

From Figure 7 we clearly see how prioritizing the old for vaccinations leads to higher (asymptomatic) infections in the spring of 2021. This reflects the fact that the young are more likely to spread the disease, since they work, so targeting vaccines to the old is the least effective strategy in terms of slowing virus diffusion. However, since the young also tend to have less severe disease progressions, both hospitalizations and deaths are lower in the spring under the old-first policy, relative to those under the young-first or the baseline all-ages-together policies.

Once Delta strikes, the advantages of the old-first vaccination strategy, in terms of avoiding

Table 4: Welfare Gains Relative to Empirical Vaccines, Empirical Mitigation, Delta

	Young First	Old First
Utilitarian Welfare	-0.08%	0.18%
Old Welfare	-1.10%	2.18%
Deaths Avoided	-51,190	104,559
GDP Gain, 2021	-0.39%	0.63%

high levels of hospitalizations and deaths, are even more stark. In part this is because the vulnerable old are almost all vaccinated, and thus even though vaccines are not fully protective, fewer old people become seriously ill, and hospital capacity is only mildly exceeded. A second reason the old-first vaccination policy performs well is that this policy delivers more infections before the emergence of Delta (the top panel) and thus a higher level of natural immunity when the Delta variant hits the U.S. Because the old-first economy is closer to herd immunity moving into the late fall, it exhibits lower infection rates relative to those of the other economies.

One feature of Figure 7 that is somewhat surprising is that the baseline vaccination profile economy exhibits an earlier and sharper Delta wave than the other two economies. The reason is that the stock of infections at the date Delta emerges, while very low in all three economies, is highest in the baseline economy, and with more initial fuel, the Delta fire grows more rapidly at first. The young-first economy has the smallest initial level of infections, and thus Delta takes longer to fully ignite in that economy, even though it leads to the largest number of eventual hospitalizations and deaths.

The welfare consequences of an old-first and a young-first policy, relative to those of the benchmark vaccination policy, are summarized in Table 4. We observe that a young-first policy becomes very costly for the old when the Delta variant emerges: a highly contagious virus strain hits a vulnerable elderly population, leading to many additional deaths. Conversely, an old-first policy proves effective in curbing fatalities in the presence of Delta, even without the reintroduction of massive lockdowns.

6 Conclusion

The development and deployment of an effective vaccine fundamentally changed the evolution of the Covid-19 pandemic in 2021. The main policy questions, given a supply of vaccines that was relatively fixed in the short run, were whom to vaccinate first, and by how much the disease progression should be slowed down by non-pharmacological interventions (such as restricting economic activities in certain sectors of the economy) as vaccinations slowly progressed.

In most countries vaccines were first given predominantly to the elderly, since this group is most likely to die from Covid-19, conditional on infection. However, since the young work and a significant share of infections occurs in the workplace, vaccinating the young first would have been an alternative that could have slowed down the pandemic even more drastically, in turn potentially necessitating costly economic mitigation policies. This paper evaluates the merit of these arguments in the epidemiological-economic model with household heterogeneity developed in Glover et al. (2021).

We find that under a utilitarian social welfare function, vaccinating the old first is preferable to vaccinating the young first, although young workers (85% of the population) sustain welfare losses, relative to a young-first policy. The endogenous adjustment of economic lockdowns reduces (but does not offset) these differences in policy preferences across the population by loosening the extent of lockdowns when the old get vaccinated first and by strengthening them when it is the young who are given vaccine priority.

We want to conclude by highlighting potential extensions of our work to relax some of the possibly strong assumptions we have made. First, in our model, even though the dynamics of the pandemic are impacted by economic activity and therefore can be partially controlled by economic mitigation policies, private economic choices (e.g., how much to work, whom to interact with) do not affect individual infection risk. Therefore, private mitigation efforts to protect one's health are by construction absent, which in turn might affect the impact of public mitigation and vaccination policies. Brotherhood et al. (2020), Eichenbaum et al. (2021), Farboodi et al. (2021), Krueger et al. (2020), Rowthorn and Toxvaerd (2020) and Toxvaerd

(2020) make these considerations the centerpiece of their analyses of the Covid-19 pandemic.[11]

Second, our analysis focused on two sources of heterogeneity, age and sector of work, we think are perhaps most salient for the Covid-19 pandemic. It thereby abstracts from the very significant regional heterogeneity in the timing and the severity of the pandemic and thus its desired control by pharmacological and non-pharmacological interventions, both across U.S. regions and across countries.[12] Even along the age dimension, it excludes school-age children and the potential effects of non-pharmacological interventions on school closures, learning outcomes and human capital accumulation.[13]

Third, the emergence of the Delta variant, together with declining immunity from first vaccinations, raises the questions of the desirability of booster shots and for whom to deploy these first, in the fall and winter of 2021 and 2022. The appearance of the Omicron variant in South Africa, and possibly already in Europe, Asia and North America, makes this analysis all the more urgent.

Such an extension calls for modeling different virus variants explicitly in the context of our (or related) EPI-Econ class of models and conducting an investigation of vaccine deployment in a world in which these vaccines are differentially effective against the various mutations. Future work is needed to investigate whether these omitted dimensions of the Covid-19 crisis reinforce or partially overturn the main conclusions of this paper—namely, that vaccines deliver large welfare benefits and that it was optimal to prioritize the elderly for their use.

References

Abel, A. B. and S. Panageas (2020): "Social Distancing, Vaccination and the Paradoxical Optimality of an Endemic Equilibrium," Working Paper 27742, National Bureau of Economic Research.

[11] This literature in turn builds on models of behavior when faced with other health risks; see, e.g., Kremer (1996) and Greenwood et al. (2019) for the AIDS epidemic.

[12] See Alon et al. (2021) for cross-country evidence on the impact of Covid-19.

[13] A very recent literature studies the impact of school closures on student human capital and future earnings, both in the short run and in the long run, once these children cohorts have entered the labor market. See, e.g., Agostinelli et al. (2020), Fuchs-Schündeln et al. (2020) or Fuchs-Schündeln et al. (2021) for papers in this line of work.

Acemoglu, D., V. Chernozhukov, I. Werning, and M. Whinston (2021): "Optimal Targeted Lockdowns in a Multi-Group SIR Model," American Economic Review: Insights, 3, 487–502.

Agostinelli, F., M. Doepke, G. Sorrenti, and F. Zilibotti (2020): "When the Great Equalizer Shuts Down: Schools, Peers, and Parents in Pandemic Times," Working Paper 28264, National Bureau of Economic Research.

Alon, T., M. Kim, D. Lagakos, and M. VanVuren (2021): "Macroeconomic Effects of COVID-19 across the World Income Distribution," Working Paper, UCSD.

Argente, D., F. Alvarez, and F. Lippi (2021): "A Simple Planning Problem for COVID-19 Lock-down, Testing, and Tracing," American Economic Review: Insights, 3, 367–382.

Berger, D., K. Herkenhoff, C. Huang, and S. Mongey (2022): "Testing and re-opening in an SEIR Model," Review of Economic Dynamics, 43, 1–21.

Bognanni, M., D. Hanley, D. Kolliner, and K. Mitman (2020): "Economics and Epidemics: Evidence from an Estimated Spatial Econ-SIR Model," Centre for Economic Policy Research, DP 13797.

Boppart, T., K. Harmenberg, J. Hassler, P. Krusell, and J. Olsson (2020): "Integrated Epi-Econ Assessment," Working Paper 28282, National Bureau of Economic Research.

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, Discussion Paper No. 175.

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-ccenarios-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, 4, 1–14.

Eichenbaum, M., S. Rebelo, and M. Trabandt (2021): "The Macroeconomics of Epidemics," Review of Financial Studies, 34, 5149–5187.

Farboodi, M., G. Jarosch, and R. Shimer (2021): "Internal and External Effects of Social Distancing in a Pandemic," Journal of Economic Theory, 196.

Fuchs-Schündeln, N., D. Krueger, A. Kurmann, E. Lale, A. Ludwig, and I. Popova (2021): "The Fiscal and Welfare Effects of Policy Responses to the Covid-19 School Closures," Working Paper 29398, National Bureau of Economic Research.

Fuchs-Schündeln, N., D. Krueger, A. Ludwig, and I. Popova (2020): "The Long-Term Distributional and Welfare Effects of COVID-19 School Closures," Working Paper 27773, National Bureau of Economic Research.

Garriga, C., R. E. Manuelli, and S. Siddhartha (2020): "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 (2021): "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic," Working Paper, University of Pennsylvania.

Gollier, C. (2021): "The Welfare Cost of Vaccine Misallocation, Delays and Nationalism," Centre for Economic Policy Research, DP 15980.

Gonzalez-Eiras, M. and D. Niepelt (2020): "On the Optimal 'Lockdown' During an Epidemic," Centre for Economic Policy Research, DP 14612.

Greenwood, J., P. Kircher, C. Santos, and M. Tertilt (2019): "An Equilibrium Model of the African HIV/AIDS Epidemic," Econometrica, 87, 1081–1113.

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," Conference Presentation.

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." Working Paper 27794, National Bureau of Economic Research.

Kremer, M. (1996): "Integrating Behavioral Choice into Epidemiological Models of AIDS," Quarterly Journal of Economics, 111, 549–573.

Krueger, D., H. Uhlig, and T. Xie (2020): "Macroeconomic Dynamics and Reallocation in an Epidemic," Centre for Economic Policy Research, DP 14607.

Makris, M. and F. Toxvaerd (2021): "Great Expectations: Social Distancing in Anticipation of Pharmaceutical Innovations," Working Paper, University of Cambridge.

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.

Rowthorn, R. and F. Toxvaerd (2020): "The Optimal Control of Infectious Diseases via Prevention and Treatment," Working Paper, University of Cambridge.

Toxvaerd, F. (2020): "Equilibrium Social Distancing," Cambridge-INET Working Paper Series No: 2020/08.

Appendix: Not for Publication

A Details of the SAFER Model

The remainder of the epidemiological SAFER model describes the transition of individuals though the health states (asymptomatic, fever-suffering, hospitalized, and recovered) once they have been infected. Its mathematical formulation is given as follows:

$$\dot{x}^{yba} = -\dot{x}^{ybs} - \left(\sigma^{yaf} + \sigma^{yar}\right) x^{yba} \tag{24}$$
$$\dot{x}^{y\ell a} = -\dot{x}^{y\ell s} - \left(\sigma^{yaf} + \sigma^{yar}\right) x^{y\ell a} \tag{25}$$
$$\dot{x}^{oa} = -\dot{x}^{os} - \left(\sigma^{oaf} + \sigma^{oar}\right) x^{oa} \tag{26}$$
$$\dot{x}^{ybf} = \sigma^{yaf} x^{yba} - \left(\sigma^{yfe} + \sigma^{yfr}\right) x^{ybf} \tag{27}$$
$$\dot{x}^{y\ell f} = \sigma^{yaf} x^{y\ell a} - \left(\sigma^{yfe} + \sigma^{yfr}\right) x^{y\ell f} \tag{28}$$
$$\dot{x}^{of} = \sigma^{oaf} x^{oa} - \left(\sigma^{ofe} + \sigma^{ofr}\right) x^{of} \tag{29}$$
$$\dot{x}^{ybe} = \sigma^{yfe} x^{ybf} - \left(\sigma^{yed} + \sigma^{yer}\right) x^{ybe} \tag{30}$$
$$\dot{x}^{y\ell e} = \sigma^{yfe} x^{y\ell f} - \left(\sigma^{yed} + \sigma^{yer}\right) x^{y\ell e} \tag{31}$$
$$\dot{x}^{oe} = \sigma^{ofe} x^{of} - \left(\sigma^{oed} + \sigma^{oer}\right) x^{oe} \tag{32}$$
$$\dot{x}^{ybr} = \sigma^{yar} x^{yba} + \sigma^{yfr} x^{ybf} + (\sigma^{yer} - \varphi) x^{ybe} \tag{33}$$
$$\dot{x}^{y\ell r} = \sigma^{yar} x^{y\ell a} + \sigma^{yfr} x^{y\ell f} + (\sigma^{yer} - \varphi) x^{y\ell e} \tag{34}$$
$$\dot{x}^{or} = \sigma^{oar} x^{oa} + \sigma^{ofr} x^{of} + (\sigma^{oer} - \varphi) x^{oe} \tag{35}$$
$$\tag{36}$$

Equations (24) to (26) give the change in the number of asymptomatic individuals. Entry into that state occurs from the inflow of newly infected individuals described in the main text. Exit from this state to the fever state occurs at rate σ^{yaf} (σ^{oaf}) for the young (old), and exit to the recovered health status occurs at rate σ^{yar} (σ^{oar}) for the young (old).

Entry into the fever state is from the asymptomatic state—see equations (27) to (29)—with exit occurring to the hospitalized state at rate σ^{yfe} and to the recovered state at rate σ^{yfr}. The old face similar transitions. Equations (30) to (32) display the transitions of those in emergency care, with entry from the fever state and exits to death and recovery. The death rate is given by $\sigma^{yed} + \varphi$, and the recovery rate is given by $\sigma^{yer} - \varphi$, where φ stems from excess death emerging from from hospital overuse. Finally, Equations (33) to (35) show the evolution of the mass of the recovered population, and the evolution of the deceased population is determined by $\dot{x}^{ybd} = (\sigma^{yed} + \varphi)x^{ybe}$, $\dot{x}^{y\ell d} = (\sigma^{yed} + \varphi)x^{y\ell e}$, and $\dot{x}^{od} = (\sigma^{oed} + \varphi)x^{oe}$.

B Parameter Values for the Benchmark Model

All epidemiological and economic parameter values used in the quantitative exercises are summarized in Tables 5 and 6.

Table 5: 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 10

This calibration implies the Spring 2020 population distribution by health status described in Table 7.

Table 6: 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
$\overline{u}$	value of life	VSL $= 10.8\times$ consumption p.c.	11.61
$\widehat{u}^f$	disutility of fever	lose 30% of baseline utility	-3.24
$\widehat{u}^e$	disutility of emergency care	lose 100% of baseline utility	-10.8
η	elasticity lux. demand to hospitalizations	CPI relative prices	-156.5
	Technology and Fiscal Policy		
μ^b	size of basic sector	55%	0.55
g	pre-COVID govt. spending	24.7% of GDP	0.247
τ^*	pre-COVID tax rate	utilitarian optimal	0.303
T^*	pre-COVID transfer	budget balance	0.223
Θ	hospital capacity	$100,000$ beds	0.000303
λ_o	impact of overuse on mortality	25% higher mortality at 200,000	825

Table 7: 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