The Distributional Effects of COVID-19 and Optimal Mitigation Policies

Sewon Hur

The Distributional Effects of COVID-19 and Optimal Mitigation Policies[*]

Sewon Hur[†]

First draft: August 2020
This draft: May 2022

Abstract

This paper develops a quantitative heterogeneous agent–life cycle model with a fully integrated epidemiological model in which economic decisions affect the spread of COVID-19 and vice versa. The calibrated model is used to study the distributional consequences and effectiveness of mitigation policies such as a stay-at-home subsidy and a stay-at-home order. First, the stay-at-home subsidy is preferred because it reduces deaths by more and output by less, leading to a larger average welfare gain that benefits all individuals. Second, Pareto-improving mitigation policies can reduce deaths by nearly 45 percent without any corresponding reduction in output relative to no public mitigation. Finally, it is possible to simultaneously improve public health and economic outcomes, suggesting that debates regarding a supposed tradeoff between economic and health objectives may be misguided.

Keywords: pandemic, coronavirus, COVID-19, mitigation, tradeoffs.

JEL Classification Codes: D62, E21, E32, E62, I14, I15.

1 Introduction

Amid the deadliest pandemic since the 1918 influenza outbreak and the largest economic contraction since the Great Depression, policymakers and intellectuals have debated a supposed tradeoff between economic and public health outcomes. On one end of the spectrum, President Trump on May 6, 2020, asked rhetorically, "Will some people be affected badly?" and responded "Yes, but we have to get our country opened and we have to get it open soon."[2] On the other end, New York Governor Cuomo on March 23, 2020, tweeted, "If it's public health versus the economy, the only choice is public health."[3] In this paper, however, I show that it is possible to simultaneously improve public health and economic outcomes, suggesting that there need not be a tradeoff between economic and health objectives.

To better understand the economic–health tradeoff (or lack thereof), I build a quantitative model that I use as a laboratory to investigate the effects of various mitigation policies. Building on the economic-epidemiological model developed by Eichenbaum et al. (2020) that allows for rich feedback between economic activities and the spread of the virus, I add three important ingredients: heterogeneity in age, heterogeneity in income and wealth, and a distinction between inside (or safe) economic activities versus outside (or risky) ones. Age heterogeneity is a necessary consideration since COVID-19 is particularly dangerous for older individuals, while mitigation policies that restrict economic activity more adversely affect working-age individuals. Heterogeneity in income and wealth is necessary in order to consider the heterogeneous effects of mitigation policies—such as stay-at-home orders—that may disproportionately harm low-wage workers, who are less likely to work from home, and low-wealth workers, who lack the resources to weather prolonged time away from work. Since mitigation policies focus on reducing economic activities that may contribute to spreading the virus, it is important to allow for inside economic activities that do not, such as working and consuming from home.

Using the model calibrated to the COVID-19 pandemic, I show that when governments only use stay-at-home orders (lockdowns) to mitigate the virus, saving more lives leads to reduced output, generating a tradeoff between lives and the economy. As an alternative, when the government pays individuals to stay at home (and not work) by offering stay-at-home subsidies, it is possible to save lives and increase output, relative to no mitigation policy. In fact, relative to no mitigation policy, a weekly subsidy of $500 can increase output

---

[2]See   https://www.cnbc.com/2020/05/05/trump-acknowledges-some-coronavirus-deaths-will-result-from-reopening.html.

[3]See https://twitter.com/nygovcuomo/status/1242264009342095361?lang=en.

by nearly 1 percent and reduce deaths by nearly 30 percent. Larger subsidies can reduce deaths by up to 45 percent without any corresponding reduction in output, relative to no mitigation policy.

How can a policy that subsidizes individuals to reduce their labor supply lead to increased output? The increase is the result of two opposing effects. The first is the *direct effect*: The subsidy provides an incentive to not work (holding fixed the severity of the pandemic), which leads to less output. This effect is small for subsidies less than $500 per week because they only change the behavior of very low-wage workers. The second is the *indirect effect*: The mitigation policy attenuates the pandemic, leading to increased economic activity. For mitigation policies with moderate subsidy amounts, the indirect effect dominates, leading to better economic and health outcomes. This is possible because the labor from the lowest-wage workers is almost exclusively "outside," implying that most of the reduction in hours contribute to mitigating the severity of the pandemic. Thus, stay-at-home subsidies change the composition of labor supply: They decrease the hours of low-wage workers (who mostly cannot work from home) and increase the hours of high-wage workers (who mostly but not entirely work from home), simultaneously increasing production and mitigating the severity of the pandemic.

To quantify these effects and analyze optimal mitigation policies, in the first part of the paper I develop a quantitative heterogeneous agent–life cycle model with a fully integrated epidemiological model in which economic activities such as outside consumption and outside labor affect the spread of COVID-19, and conversely, the virus affects economic decisions. The model also features incomplete markets, endogenous labor with the option to work from home, and hospital capacity constraints.

The model has two main sources of externalities. The first is static and is generated by *currently infected* individuals who do not take into account how their economic activities may transmit the virus to susceptible individuals. The second externality is dynamic and is generated by *currently susceptible* individuals who understand how their economic activities affect their own probability of getting infected, but do not take into account how those activities could transmit the virus if they become infected. This latter externality is exacerbated by the fact that COVID-19 fatality risk varies across age groups.

In the second part of the paper, I describe the model's calibration. The model is solved at a biweekly frequency to study the progression of the disease at a high frequency. The model's economic parameters are calibrated to match both aggregate and distributional features of the US economy before the pandemic and the model's epidemiological and clinical parameters

are set to match features of COVID-19, such as estimates for the basic reproduction number, age-specific fatality rates, and the time series of COVID-19 deaths in the United States. I also show that the model matches key features of the data, both targeted and non-targeted, and generates time series for relevant aggregate variables such as outside consumption and labor that match the data reasonably well.

In the third part of the paper, I use the calibrated model to quantify the heterogeneous responses to the pandemic across individuals.

The first main finding is that, **in the absence of mitigation policies, private mitigation by individuals is heterogeneous.** Individuals voluntarily reduce their outside consumption and hours worked to reduce their probability of infection. While this is a common feature in economic-epidemiological models, such as Eichenbaum et al. (2020), the rich heterogeneity in my model allows for additional new insights. All else equal, private mitigation is stronger for older individuals, who face higher death rates if infected; for higher-wage workers, who are more likely to work from home; and for wealthier individuals, who can afford to sustain prolonged time away from work. This highlights the dynamic externality explained above: Young low wage–low wealth (susceptible) workers engage in too much economic activity, relative to the social optimum, leading to higher infections and deaths in the aggregate. Additionally, low-wage workers' reduction of outside consumption and hours by less than their high-wage counterparts is qualitatively consistent with the fact that higher income locations had larger declines in spending and mobility than lower income locations (Chetty et al. 2020). It is also consistent with high-income individuals reducing their outside labor more than their low-income counterparts, documented using individual-level survey data provided by Bick et al. (2020).

These externalities give rise to the possibility of welfare-improving government interventions. Thus, I also use the calibrated model to compare and contrast the effects of a stay-at-home order (lockdown) that imposes a cap on outside hours worked—resembling the various stay-at-home and shelter-in-place orders implemented by local and state governments in response to the pandemic—and a stay-at-home subsidy that provides a weekly subsidy for individuals who do not work.

The second main finding is that **the stay-at-home subsidy is superior to the stay-at-home order**. Relative to the lockdown policy alone, the subsidy policy alone delivers a higher average welfare gain while reducing deaths by more and output by less. In the case of the lockdown, older individuals experience a welfare gain because of the reduced infection and death probability, but these gains are mostly offset by the welfare losses of

low-wage workers, who face a large decline in their income. In contrast, the stay-at-home subsidy can be Pareto improving. This contrast arises because while both policies result in reduced economic activities of young, low-wage/low-wealth workers, the subsidy policy provides the incentives to do so and the lockdown does not. Neither policy has a direct effect on the labor supply of high-wage workers—who choose to work mostly from home during the pandemic—though they do benefit from a milder pandemic as a result of either mitigation policy.

In the final part of the paper, I study the optimal configuration of these policies by varying the stay-at-home subsidy amount from $0 to $1,200 per week, the subsidy duration from 4 to 14 months, and the speed at which the subsidy phases out, with and without lockdowns of different intensities. In these counterfactual exercises, I assume that the subsidy is fully funded by a tax on consumption. I utilize high performance computing to solve for over 40,000 transition paths, including sensitivity analyses.

The third main finding is that **it is possible to simultaneously improve public health and economic outcomes**. The output-maximizing policy, which involves a weekly subsidy of $500, a duration of 10 months, and no lockdown, *reduces* deaths by 27 percent and *increases* two-year output by nearly 1 percentage point, compared with no public mitigation. Furthermore, this policy is Pareto improving. Larger subsidies can reduce deaths by up to 45 percent without a corresponding reduction in output, relative to no public mitigation. Finally, I find that the best Pareto improving policy—which involves a weekly subsidy of $700, a duration of 10 months, and no lockdown—reduces deaths by 36 percent, while also slightly increasing output, relative to no public mitigation.

These mitigation policies are Pareto improving even though the subsidies are fully funded with a consumption tax. Low-wage workers who choose to stay at home and accept the subsidy are better off by revealed preference. Higher-wage workers and retired individuals, who are not directly affected by the subsidy, benefit from the resulting attenuated pandemic and thus are willing to pay the higher consumption tax that is used to fund the subsidy. Furthermore, I show that a similar subsidy-and-tax policy is not welfare improving in the absence of a pandemic. This suggests that the externalities generated during the pandemic are essential for allowing for Pareto-improving policies.

Related literature. The epidemiological part of the model borrows from the SIR model of disease transmission, originally developed by Kermack and McKendrick (1927). Atkeson (2020) was one of the first papers to use the SIR model to study COVID-19 in an economics

4

context.[4] This literature is now very large. Alvarez et al. (2020), Eichenbaum et al. (2020), Farboodi et al. (2020), and Jones et al. (2020) study optimal mitigation in SIR models extended with lockdowns, economic-epidemiological feedback, social distancing, and work from home with learning-by-doing, respectively. Bodenstein et al. (2020), Baqaee et al. (2020), and Krueger et al. (2020) study the SIR model with multiple sectors. Birinci et al. (2020), Garibaldi et al. (2020), and Kapicka and Rupert (2020) incorporate search and matching frictions into the SIR framework, while Berger et al. (2020), Chari et al. (2020), and Piguillem and Shi (2020) extend the SIR model to focus on testing and quarantine. Argente et al. (2020) and Azzimonti et al. (2020) enrich the SIR model with city structure and contact networks, respectively. Bognanni et al. (2020) develop a SIR model with multiple regions and estimate it on daily county-level US data and Fernández-Villaverde and Jones (2020) estimate a SIR model for many cities, states, and countries. Aum et al. (2020) study the effects of lockdowns in a model with heterogeneous age, skill, and occupation choice and Brotherhood et al. (2020b) study age-specific testing and quarantine policies.[5]

This paper is most related to Bairoliya and Imrohoroglu (2020), Glover et al. (2020), Kaplan et al. (2020), and Nakajima (2020). Bairoliya and Imrohoroglu (2020) study targeted lockdowns in a life-cycle model with heterogeneity across age, health, income, and wealth. They find that lockdowns targeted toward individuals with greater risk to COVID-19 can improve economic outcomes relative to lockdowns that are random. To my knowledge, Nakajima (2020) is the only other paper that uses a quantitative economic-epidemiological model that features heterogeneity across age, income, and wealth. Nakajima (2020) focuses on evaluating US policies, while my paper focuses on characterizing optimal mitigation policies. Glover et al. (2020) study optimal mitigation policies in a model with three types of agents: retirees, young workers in the essential sector, and young workers in the non-essential sector. Relative to their work, my paper features heterogeneity across not only age, but also income and wealth, and complements Glover et al. (2020) by analyzing mitigation policies specifically targeting the heterogeneous behavior of these subgroups. Kaplan et al. (2020) do not model heterogeneity by age, but include heterogeneity across income, wealth, sector, and occupation. Like my paper, they focus on policies that improve not only average outcomes but also account for the distributional consequences of mitigation policies. Relative to Kaplan et al. (2020), my paper addresses the externalities that are generated by the

---

[4]The SIR model has been used in economics to study other diseases, such as AIDS (Greenwood et al. 2019; Kremer 1996).

[5]Given the rapidly expanding literature, this is likely not an exhaustive list. I refer the reader to Brodeur et al. (2020) and Hur and Jenuwine (2020) for reviews of this literature.

differential effects of COVID-19 by age. It also explicitly accounts for the value of life so that individual welfare changes directly reflect economic as well as expected health outcomes and proposes policies that can improve aggregate outcomes in a Pareto improving sense.

Numerous papers have studied the tradeoff between health and economic outcomes, such as Aspri et al. (2021), Assenza et al. (2020), Hall et al. (2020), and Mendoza et al. (2020). Many papers have also questioned the tradeoff between economic and health outcomes as this paper does. For example, Acemoglu et al. (2020) and Fajgelbaum et al. (2021) show that the tradeoff can be improved when implementing lockdowns that target heterogeneous groups and localities, respectively. As another example, Aum et al. (2020) show that implementing South Korea's testing and quarantine policies in the United Kingdom would have reduced both deaths and output losses there. Relative to these papers, my paper focuses on suggesting Pareto-improving policies in a model with rich heterogeneity that improve both health and economic objectives.

This paper is structured as follows. The next section presents the model. Section 3 describes the calibration of the model's economic and epidemiological parameters and discusses the model's fit. In Section 4, the calibrated model is used to investigate the role of private mitigation and the welfare consequences of the pandemic and mitigation policies. Section 5 discusses the properties of optimal policies. Finally, section 6 concludes.

2 Model

This section presents a model economy used to quantitatively analyze the welfare consequences of COVID-19 and to run policy counterfactuals. The setting combines a heterogeneous-agent overlapping-generations model with an economic-epidemiological model that resembles that used in Eichenbaum et al. (2020). The economy is inhabited by overlapping generations of stochastically aging individuals. Time is discrete and indexed by $t = 0, ..., \infty$. Workers face idiosyncratic efficiency shocks and borrowing constraints within an incomplete market setting. I now describe the model in more detail.

2.1 Individuals

Epidemiological block.   An individual's health status is given by $h \in \{S, I, R, D\}$: *Susceptible* agents may contract the virus; *infected* agents have contracted the virus and may pass it onto others; and agents that exit the infection can either *recover* or *die*. Let $H \equiv \{S, I, R\}$ denote the health states of living individuals. Recovered agents are assumed to be immune

from further infection.[6] The transitions between health states build on the widely used SIR model, originally developed by Kermack and McKendrick (1927).[7] Susceptible individuals get infected with probability π_{It}, which depends on individual outside consumption and outside labor (c_o, ℓ_o) and the aggregate measure of infected individuals (μ_{It}) and their outside consumption and outside labor (C_{It}^o, L_{It}^o).[8] Formally,

$$\pi_{It}(c_o, \ell_o; Z_t) = \beta_c c_o C_{It}^o + \beta_\ell \ell_o L_{It}^o + (\beta_e + \epsilon_t) \mu_{It}, \tag{1}$$

where ϵ_t captures time-varying transmissibility (e.g. seasonal factors) and $Z_t \equiv \{\mu_{It}, C_{It}^o, L_{It}^o, \epsilon_t\}$. This framework allows the virus to be contracted from consumption-related activities, labor-related activities, and from other settings. It also allows a feedback between disease progression and economic activities as in Eichenbaum et al. (2020), Glover et al. (2020), and Jones et al. (2020). Relative to these papers, however, the richer heterogeneity in this model allows for infection probabilities to be very different across the distribution.

Infected individuals exit the infection with probability π_{Xt}, and upon exit, they recover with probability $1 - \delta_{jt}(\mu_{It})$ and die with probability $\delta_{jt}(\mu_{It})$, where j is the individual's age. The fatality rate depends on the individual's age and on the aggregate measure of infected individuals, reflecting hospital capacity constraints. Finally, susceptible individuals may transition directly to the recovered state with probability π_{jRt}, for example, by a vaccine when it becomes available. Then the transition matrix between health states is given by

$$\Pi_{jhh't}(c_0, \ell_0; Z_t) = \tag{2}$$

	S	I	R	D
S	$1 - \pi_{It}(c_0, \ell_0; Z_t) - \pi_{jRt}$	$\pi_{It}(c_0, \ell_0; Z_t)$	π_{jRt}	0
I	0	$1 - \pi_{Xt}$	$\pi_{Xt}(1 - \delta_{jt}(Z_t))$	$\pi_{Xt}\delta_{jt}(Z_t)$
R	0	0	1	0
D	0	0	0	1

.

[6]One could easily extend the model to have shorter durations of immunity.

[7]A popular variant of the SIR model is the SEIR model, which adds a category of individuals that have been *exposed* to the virus, and may or may not be infectious without symptoms. This distinction is more relevant for studying disease dynamics at an even higher frequency than in this paper, and particularly for those that study testing and quarantine such as Berger et al. (2020).

[8]The distinction between inside and outside activities is based on whether or not there is infection risk and not based on physical location.

Economic block. Individuals of age $j \in J \equiv \{1, 2, .., \overline{J}\}$ face conditional aging probabilities given by $\{\psi_j\}$, with mandatory retirement at age $j = J_R$.[9] Workers $(j < J_R)$ face uninsurable wage risk: Each period, workers receive idiosyncratic efficiency shocks $\varepsilon \in E$, which follow a Markov process, with transition matrix Γ. Their labor income is given by $w_t \eta_{jh} \varepsilon \ell$, where w_t is the efficiency wage, η_{jh} is the health- and age-profile of efficiency units, and ℓ is total hours worked. Workers may choose to work up to a fraction $\bar{\theta}_j(\varepsilon)$ of their labor hours from home, where $\bar{\theta}_j(\varepsilon)$ is allowed to vary by age and efficiency. Retirees $(j \geq J_R)$ are assumed to receive a fixed income of s each period.[10] Retirement income is funded by taxes on labor income, $\tau_{\ell t}$. Individuals can accumulate non-contingent assets k, which deliver a net return of r_t.

Mitigation policies. Individuals take as given a sequence of taxes, subsidies, and other restrictions on quantities that are designed to mitigate the severity of the pandemic.

1. Stay-at-home subsidies: Individuals receive a subsidy of T_t each period in which they supply zero hours of total labor.

2. Stay-at-home order (lockdown): Lockdowns restrict the quantity of outside labor to at most $\bar{\ell}_{ot}$.

3. Taxes: In the baseline model, I consider taxes on consumption τ_{ct}. These taxes serve two purposes: a) They can be used to fund the stay-at-home subsidies and b) they reduce economic activities that contribute to the spread of the virus. In extensions of the model, I also consider taxes on labor income and Pigouvian taxes that directly target outside consumption.

Individuals' recursive problem. Given the sequence of prices $\{r_t\}$, consumption taxes $\{\tau_{ct}\}$, and aggregate states $\{Z_t\}$, a retiree with age $j \geq J_R$, wealth k, and health h in period

---

[9]Given that the model will be used to analyze disease progression at a high (bi-weekly) frequency, the assumption of stochastic aging greatly reduces the state space and computational burden.

[10]This can readily be extended to depend on lifetime earnings as in Hur (2018).

t chooses inside and outside consumption $\{c_i, c_o\}$ and savings k' to solve:

$$v_{jt}^R(k, h) = \max_{c_i, c_o, k' \geq 0} \; u(c_i, c_o) + \bar{u} \tag{3}$$
$$+ \beta \psi_j \sum_{h' \in H} \Pi_{jhh't}(c_o, 0; Z_t) v_{j+1,t+1}^R(k', h')$$
$$+ \beta(1 - \psi_j) \sum_{h' \in H} \Pi_{jhh't}(c_o, 0; Z_t) v_{j,t+1}^R(k', h')$$
$$\text{s.t. } (1 + \tau_{ct})c + k' \leq s + k(1 + r_t)$$

where β is the time discount factor, $u(c_i, c_o)$ is the utility derived from inside and outside consumption, $c = c_i + c_o$ is total consumption, and $\bar{u}$ governs the flow value of being alive. Solving this yields retiree policy functions $\{c_{ijt}^R(k, h), c_{ojt}^R(k, h), k_{jt}^{R'}(k, h)\}_{j \geq J_R}$ for inside and outside consumption and savings, respectively. I normalize the value of death to zero and set $v_{\bar{J}+1,t}^R = 0$. Thus, agents in the last stage of life $(j = \bar{J})$ may die due to stochastic aging (with probability ψ_j) and, if infected, because of the virus (with probability Π_{jIDt}).

Given the sequence of prices $\{w_t, r_t\}$, fiscal and mitigation policies $\{\tau_{ct}, \tau_{\ell t}, T_t, \bar{\ell}_{ot}\}$, and aggregate states $\{Z_t\}$, a worker with age $j < J_R$, wealth k, efficiency ε, and health h in period t chooses inside and outside consumption $\{c_i, c_o\}$, inside and outside labor $\{\ell_i, \ell_o\}$, and savings k' to solve:

$$v_{jt}(k, \varepsilon, h) = \max_{c_i, c_o, \ell_i, \ell_o, k' \geq 0} \; u(c_i, c_o) - g(\ell) + \bar{u} \tag{4}$$
$$+ \beta \psi_j \sum_{\varepsilon' \in E} \sum_{h' \in H} \Gamma_{\varepsilon\varepsilon'} \Pi_{jhh't}(c_o, \ell_o; Z_t) v_{j+1,t+1}(k', \varepsilon', h')$$
$$+ \beta(1 - \psi_j) \sum_{\varepsilon' \in E} \sum_{h' \in H} \Gamma_{\varepsilon\varepsilon'} \Pi_{jhh't}(c_o, \ell_o; Z_t) v_{j,t+1}(k', \varepsilon', h')$$
$$\text{s.t. } (1 + \tau_{ct})c + k' \leq w_t \eta_j^h (1 - \tau_{\ell t}) \varepsilon \ell + k(1 + r_t) + T_t(\ell)$$
$$\ell_i \leq \bar{\theta}_j(\varepsilon) \ell, \;\; \ell_o \leq \bar{\ell}_{ot}$$

where $\ell = \ell_i + \ell_o$ is total labor, $g(\ell)$ is the disutility of labor, and $v_{jt}(k, \varepsilon, h) = v_{jt}^R(k, h)$ for $j \geq J_R$ and $\varepsilon \in E$. Solving this yields worker policy functions $\{c_{ijt}(k, \varepsilon, h), c_{ojt}(k, \varepsilon, h), \ell_{ijt}(k, \varepsilon, h), \ell_{ojt}(k, \varepsilon, h), k'_{jt}(k, \varepsilon, h)\}_{j < J_R}$ for inside and outside consumption, inside and outside labor, and savings, respectively. Additionally, let $c_{ijt}(k, \varepsilon, h) = c_{ijt}^R(k, h)$, $c_{ojt}(k, \varepsilon, h) = c_{ojt}^R(k, h)$, and $k'_{jt}(k, \varepsilon, h) = k_{jt}^{R'}(k, h)$ for $j \geq J_R$ and $\varepsilon \in E$.

2.2 Production

A representative firm hires labor (L_{ft}) and capital (K_{ft}) to produce according to

$$Y_{ft} = K_{ft}^{\alpha} L_{ft}^{1-\alpha}. \tag{5}$$

Taking prices as given, the firm solves

$$\max_{L_{ft}, K_{ft}} Y_{ft} - w_t L_{ft} - (r_t + \delta) K_{ft}, \tag{6}$$

where δ is the deprecation rate of capital. Optimality conditions are given by

$$w_t = (1 - \alpha) K_{ft}^{\alpha} L_{ft}^{-\alpha}, \tag{7}$$

$$r_t = \alpha K_{ft}^{\alpha-1} L_{ft}^{1-\alpha} - \delta. \tag{8}$$

2.3 Equilibrium

We are ultimately interested in studying disease dynamics and mitigation policies along a transition path. However, because most of the model parameters are calibrated to an initial pre-pandemic steady state, it is useful to first define a stationary equilibrium in which $\mu_I = 0$. In this case, aggregate outside consumption and outside labor of infected individuals is trivially zero. Furthermore, if we set the time-varying transmissibility parameter $\epsilon_t = 0$, then $Z = (0, 0, 0, 0)$ and Π is the identity matrix (with no vaccine). I set aside all mitigation policies by assuming that $\tau_c = 0$ and $T = 0$ and setting $\bar{\ell}_o$ sufficiently large such that it is not binding for any individual. Finally, retirement income is financed by labor income taxes, and accidental bequests from death are distributed to newborns each period. Define the state space over wealth, efficiency, and health as $X = K \times E \times H$ and let a σ-algebra over X be defined by the Borel sets, $\mathcal{B}$, on X.

Definition. A *stationary recursive competitive equilibrium*, given fiscal policies $\{\tau_\ell, s\}$, is a set of value functions $\{v_j\}_{j \in J}$, policy functions $\{c_{ij}, c_{oj}, \ell_{ij}, \ell_{oj}, k'_j\}_{j \in J}$, prices $\{w, r\}$, producer plans $\{Y_f, L_f, K_f\}$, the distribution of newborns ω, and invariant measures $\{\mu_j\}_{j \in J}$ such that:

1. Given prices and fiscal policies, retirees and workers solve (3) and (4), respectively.

2. Given prices, firms solve (6).

3. Markets clear:

   (a) $Y_f = \int_X \sum_{j \in J} (c_{ij}(k, \varepsilon, h) + c_{oj}(k, \varepsilon, h) + \delta k) \, d\mu_j(k, \varepsilon, h)$,

(b) $L_f = \int_X \sum\limits_{j<J_R} \eta_{jh}\varepsilon \left(\ell_{ij}(k,\varepsilon,h) + \ell_{oj}(k,\varepsilon,h)\right) d\mu_j(k,\varepsilon,h),$

(c) $K_f = \int_X \sum\limits_{j\in J} k d\mu_j(k,\varepsilon,h).$

4. The government budget constraint holds:

$$\tau_\ell w \int_X \sum_{j<J_R} \eta_{jh}\varepsilon \left(\ell_{ij}(k,\varepsilon,h) + \ell_{oj}(k,\varepsilon,h)\right) d\mu_j(k,\varepsilon,h) = s \int_X \sum_{j\geq J_R} d\mu_j(k,\varepsilon,h). \qquad (9)$$

5. For any subset $(\mathcal{K},\mathcal{E},\mathcal{H}) \in \mathcal{B}$, the invariant measure μ_j satisfies, for $j > 1$,

$$\mu_j(\mathcal{K},\mathcal{E},\mathcal{H}) = \int_X \psi_{j-1} \mathbb{1}_{\left\{k'_{j-1}(k,\varepsilon,h)\in\mathcal{K}\right\}} \sum_{\varepsilon'\in\mathcal{E}} \sum_{h'\in\mathcal{H}} \Gamma_{\varepsilon\varepsilon'}\Pi_{jhh'} d\mu_{j-1}(k,\varepsilon,h) \qquad (10)$$

$$+ \int_X (1-\psi_j)\mathbb{1}_{\left\{k'_j(k,\varepsilon,h)\in\mathcal{K}\right\}} \sum_{\varepsilon'\in\mathcal{E}} \sum_{h'\in\mathcal{H}} \Gamma_{\varepsilon\varepsilon'}\Pi_{jhh'} d\mu_j(k,\varepsilon,h)$$

and

$$\mu_1(\mathcal{K},\mathcal{E},\mathcal{H}) = \int_X (1-\psi_1)\mathbb{1}_{\left\{k'_1(k,\varepsilon,h)\in\mathcal{K}\right\}} \sum_{\varepsilon'\in\mathcal{E}} \sum_{h'\in\mathcal{H}} \Gamma_{\varepsilon\varepsilon'}\Pi_{jhh'} d\mu_1(k,\varepsilon,h) + \omega(\mathcal{K},\mathcal{E},\mathcal{H}). \quad (11)$$

6. The newborn distribution satisfies:

$$\int_X k d\omega(k,\varepsilon,h) = \int_X \psi_{\bar{J}} k'_{\bar{J}}(k,\varepsilon,h) d\mu_{\bar{J}}(k,\varepsilon,h). \qquad (12)$$

Note that I assume that inside and outside consumption and investment use the same good. Similarly, note that inside and outside labor are perfect substitutes in production. I discuss the implications of these and other assumptions in Section 5.

3 Calibration

In this section, I begin by calibrating some of the model's parameters to the pre-pandemic steady state and discuss how other parameters that require solving for the transition path are set. I then discuss the model's fit by examining both targeted and non-targeted moments in the steady state and during the course of the pandemic. The parameters are summarized in Tables 1 and 2. See Appendix C for details regarding the computation and calibration strategy.

3.1 Calibration of the pre-pandemic steady state

Environment and demographics. A period in the model is two weeks. The aggregate measure of individuals in the steady state economy is normalized to one. The number of age cohorts, J, is set to 3, so that $j = 1$ corresponds to ages 25–44 (young), $j = 2$ corresponds to ages 45–64 (middle), and $j = J_R = \bar{J} = 3$ corresponds to ages 65+ (retired). The worker aging probability $\psi_1 = \psi_2$ is set so that workers spend, on average, 20 years in each age cohort. The aging (death) probability of retired individuals ψ_3 is set so that the retired account for 20 percent of the 25+ population. The wealth of deceased individuals is given to a fraction of newborn individuals each period. Specifically, 85 percent of individuals are born with zero wealth, whereas 15 percent of individuals are endowed with 28 times annual per capita consumption.[11]

Preferences. The utility function is assumed to take the form:

$$u(c_i, c_o) = \frac{\left(c_i^{\gamma} c_o^{1-\gamma}\right)^{1-\sigma}}{1 - \sigma}, \tag{13}$$

which exhibits constant relative risk aversion over a Cobb-Douglas aggregation of inside and outside consumption. There are two points worth discussing. First, the distinction between inside and outside consumption in the model is purely based on whether there is risk of infection. There is not such a clear distinction in the data. In order to pin down γ, I assume that household expenditures on goods and housing and utility services represent expenditures that are relatively safe, which constitute 51 percent of total household expenditures on goods and services (2019, *Bureau of Economic Analysis*). Accordingly, γ is set to 0.51 to match this share. I then show in Appendix D that the main findings of the paper are robust to alternative values of this parameter. Second, I am assuming a unitary elasticity of substitution between inside and outside consumption.[12] I show later in this section that the model generates

---

[11] This is based on the fact that 85 percent of households whose heads are between the ages of 21 and 25 had a cumulative net worth of zero in 2016 (*Survey of Consumer Finances*). The calibrated value of the endowment is rather large. This issue could be addressed by increasing the number of retired cohorts so that retired individuals draw down more wealth before dying. In sensitivity analysis (Appendix D), I show that the main results are robust to alternatively assuming that a deceased household is replaced by a newborn with the same efficiency and wealth.

[12] A unitary elasticity is somewhat consistent with estimates for the elasticity of substitution between market and home goods, which range between 0.8 (Duernecker and Herrendorf 2018) and 2.2 (Dotsey et al. 2014). However, these estimates are only partially informative since many market goods can be purchased online and consumed safely at home (e.g. consumer electronics). Boppart et al. (2020) also use a unitary elasticity of substitution between inside and outside consumption.

reasonable time series of total and outside consumption in response to the pandemic.

The disutility from labor is given by:

$$g(\ell) = \varphi \frac{\ell^{1+\nu}}{1+\nu} + \mathbb{1}_{\{\ell=0\}} \tilde{u}, \tag{14}$$

where $\tilde{u}$ represents the disutility from not working (e.g. administrative costs, stigma, or any other costs not modeled explicitly here).[13] I set $\tilde{u}$ so that not working is associated with a 5 percent reduction in the flow value of life. This generates a 15 percent reduction in employment during the transition path described in Section 3.2, matching the decline in the data (2020, *Bureau of Labor Statistics*).[14] The parameter that governs the disutility from labor, φ, is set so that the model generates an average of 34.4 hours worked per week (2019, *Bureau of Economic Analysis*).

The time discount factor β is chosen so that the model replicates the US net-worth-to-GDP ratio (2010–2019, *Board of Governors*). Finally, I set risk aversion, σ, to be 2, and the Frisch elasticity, $1/\nu$, to be 0.5—commonly used values in the literature (for example, see Heathcote et al. 2014 and Storesletten et al. 2004 for risk aversion and Chetty et al. 2011 and Kaplan et al. 2018 for the Frisch elasticity).

Technology and income. The capital elasticity in the production function, α, is set to match the aggregate capital income share of 0.36. The labor income tax τ_ℓ and retirement income s are chosen so that retirement income is 30 percent of average labor earnings in the model and the government budget constraint is satisfied. The depreciation rate of capital, δ, is set at an annualized rate of 5 percent per year, a standard value used in the literature (for example, see Kehoe et al. 2018 and Chari et al. 2007). The stay-at-home subsidy and consumption tax are set to zero in the pre-pandemic steady state.

The age-profile of efficiency units, η_{jS}, is normalized to one for susceptible young workers and susceptible middle-age workers are assumed to be 35 percent more efficient, to match the wage ratio of middle-age to young workers in the data (2014, *Panel Survey of Income Dynamics*). I assume that the efficiencies of recovered individuals are the same as those of susceptible individuals, $\eta_{jR} = \eta_{jS}$.[15] The fraction of labor that can be done from home,

---

[13]The fixed cost operates similarly to other settings with extensive margins in labor supply (e.g. French and Jones 2011) and stock market participation (e.g. Allen and Gale 1994).

[14]Sensitivity analysis (Appendix D) demonstrates that the main results are robust to a lower utility cost, which generates a higher reduction in employment that is closer to the 19 percent reduction reported by Bick et al. (2020) based on survey data.

[15]It is too early to conclude about the potentially long-lasting consequences of COVID-19. That said, these assumptions can easily be modified if evidence dictates.

Table 1: Calibration of economic parameters

Parameters	Values	Targets / Source
Discount factor, annualized, β	0.97	Wealth-to-GDP: 4.8
Risk aversion, σ	2	Heathcote et al. (2014)
Inside consumption share, γ	0.51	Expenditure share
Disutility from labor, φ	20.55	Average weekly hours: 34.4
Frisch elasticity, $1/\nu$	0.50	Chetty et al. (2011)
Death probability, annualized, ψ_3	0.10	65+ share of population 25+: 0.2
Aging probability, annualized, $\psi_1 = \psi_2$	0.05	Expected duration: 20 years
Efficiency units, $\eta_{1S} = \eta_{1R}$	1.00	Wage ratio of age 45-64 workers
$\eta_{2S} = \eta_{2R}$	1.35	to age 25-44 workers
Factor elasticity, α	0.36	Capital share
Capital depreciation, annualized, δ	0.05	Kehoe et al. (2018)
Retirement income, s	1.02	30% of average earnings per worker
Labor income tax, τ_ℓ	0.07	Government budget constraint
Persistence, annual, ρ_ε	0.94	Author estimates
Standard deviation, annual, σ_υ	0.19	Author estimates

$\bar{\theta}_j(\varepsilon)$, is set to match the average share of jobs that can be done from home by occupations grouped into five wage bins, computed based on Dingel and Neiman (2020).[16] Thus, $\bar{\theta}_j(\varepsilon)$ is set to 0.03 for individuals in the bottom 20 percent of the wage distribution, 0.21 for the second quintile, 0.32 for the third quintile, 0.47 for the fourth quintile, and 0.66 for those in the top 20 percent of the wage distribution, where the wage is defined as $w\eta_{jS}\varepsilon$.

The labor efficiency shocks ε are assumed to follow an order-one autoregressive process as follows:

$$\log \varepsilon_t = \rho_\varepsilon \log \varepsilon_{t-1} + \upsilon_t, \ \upsilon_t \sim N\left(0, \sigma_\upsilon^2\right). \tag{15}$$

This process is estimated using annual wages constructed from the PSID to find a persistence of $\rho_\varepsilon = 0.94$ and a standard deviation of $\sigma_\upsilon = 0.19$.[17] These parameters are then converted to a higher frequency, following Krueger et al. (2016). The process is approximated with a seven-state Markov process using the Rouwenhurst procedure described in Kopecky and Suen (2010).

---

[16]See Appendix A.1 for details on data construction.

[17]The wages are constructed similarly to Floden and Lindé (2001) and the sample selection and estimation procedures closely follow Krueger et al. (2016). See Appendix A.2 for details on the wage process estimation.

Clinical and epidemiological parameters. The infection exit rate, π_X, is set to 14/18 so that the expected duration of the infection is 18 days, as in Atkeson (2020) and Eichenbaum et al. (2020). For the unconstrained case fatality rates, I use data from South Korea's Ministry of Health and Welfare (accessed August 4, 2020) to compute a fatality rate of 8.47 percent for ages 65–84, 0.85 percent for ages 45–64, and 0.08 percent for ages 25–44. I use South Korean data because testing has been abundant since the outbreak began[18], the in-sample peak in infections was early enough that case fatality rates are not biased due to lags in deaths, and hospitals were not overwhelmed, as the number of active cases never exceeded 0.015 percent of the population.[19]

Next, I discuss the hospital capacity constraints and how they affect death rates. Following Piguillem and Shi (2020), I use the functional form

$$\delta_j(\mu_I) = \delta_j^u \min\left\{1, \frac{\kappa}{\mu_I}\right\} + \delta_j^c \max\left\{0, 1 - \frac{\kappa}{\mu_I}\right\} \tag{16}$$

where δ_j^u and δ_j^c denote the unconstrained and untreated death rates, respectively, and κ denotes the measure of infected individuals that can be treated without the constraint binding. According to the American Hospital Association, there are roughly 924,000 hospital beds in the US, corresponding to 0.28 percent of the population.[20] Since not all infected cases require hospitalization, I use a generous capacity constraint, κ, of 1 percent. The unconstrained death rates, δ_j^u, are set to match those documented for South Korea, and the untreated death rates are set as $\delta_j^c = 2\delta_j^u$, following Piguillem and Shi (2020). I later show that infections peak at 1.4 percent, implying a 27.5 percent increase in fatality rates at the peak (e.g. the middle-age fatality rate increases from 0.85 percent to 1.08 percent).[21]

There is some uncertainty regarding the basic reproduction number (R_0), which corresponds to the number of people to whom the average infected person passes the disease absent mitigation efforts. Most estimates range between 2.2 and 3.1 (see for example, Wang et al. 2020 and Fauci et al. 2020), so I use a conservative estimate of 2.2.[22] By aggregating

---

[18]For example, see https://www.bloomberg.com/news/articles/2020-04-18/seoul-s-full-cafes-apple-store-lines-show-mass-testing-success. Aum et al. (2020) also discuss the success of early testing and tracing efforts in South Korea.

[19]Prior to the second wave in December of 2020, active infection cases in South Korea had peaked at 7,362 on March 11, 2020, according to https://www.worldometers.info/coronavirus/country/south-korea/.

[20]See https://www.aha.org/statistics/fast-facts-us-hospitals.

[21]In sensitivity analysis (Appendix D), I show that the main findings are robust to assuming that there is no capacity constraint.

[22]It is worth noting that more recent Delta and Omicron variants are more infectious than the original strain. In sensitivity analysis (Appendix D), I show that the main results are robust to a higher value for the basic reproduction number.

equation (1) over the measure of susceptible agents and applying the law of large numbers, I obtain the expression for new infections in a given period:

$$T_t = \beta_c C_{St}^o C_{It}^o + \beta_\ell L_{St}^o L_{It}^o + (\beta_e + \epsilon_t) \mu_{St} \mu_{It}, \tag{17}$$

where C_{ht}^o and L_{ht}^o denote aggregate outside consumption and labor, respectively, of individuals with health h in period t. In the pre-pandemic steady state, workers are indifferent between working outside or working from home. Thus, I assume that all steady state work is done outside, which can be obtained by introducing an arbitrarily small difference in either efficiency or preference in favor of working outside. By substituting $L_S/\mu_S = L_I/\mu_I$ and $C_S^o/\mu_S = C_I^o/\mu_I$, taking $\mu_S \to 1$, and setting $\epsilon_t = 0$, the basic reproduction number is given by

$$R_0 = \frac{\beta_c(C_S^o)^2 + \beta_\ell(L_S^o)^2 + \beta_e}{\pi_X}. \tag{18}$$

Thus, given values for the basic reproduction number, R_0, the exit rate, π_X, the steady state values for aggregate outside consumption and labor, C_S^o and L_S^o, we need to assign values to the fractions of new infections occurring through consumption activities, work activities, and other settings to pin down the values for β_c, β_ℓ, and β_e. Evidence on how COVID-19 is transmitted is thus far limited, but in the case of other infectious diseases, Ferguson et al. (2006) report that 70 percent of transmissions occur outside of the household. In another study that investigates the transmission channels of infectious diseases, Mossong et al. (2008) find that 35 percent of high-intensity contacts occur in workplaces and schools. Roughly based on these studies, I assume that one-third of initial transmission occurs through consumption activities, one-third through labor activities, and one-third through other channels.

I set the flow value of life $\bar{u}$ so that the model generates an average value of statistical life (VSL) of $6.9 million in 2019—the value used by U.S. Federal Emergency Management Agency (2019)—corresponding to 4,096 times biweekly consumption per capita in 2019.[23] See Appendix C.1.1 for the derivation of VSL in the model.

Next, I discuss how the efficiency units change when an individual gets infected. It is reasonable to expect that those with no symptoms would suffer little, if any, efficiency loss, whereas those that experience very severe symptoms would suffer something close to a 100 percent efficiency loss. Without sufficient evidence regarding how a COVID-19 infection affects labor efficiency, I assume that infected individuals suffer a 50 percent loss in efficiency.

---

[23]As a robustness check, I use a higher VSL of $7.4 million in 2006—the value recommended by U.S. Environmental Protection Agency (2020)—or 6,226 times biweekly consumption per capita (see Appendix D). The main results of the paper are robust to this higher value.

Table 2: Clinical and epidemiological parameters

Parameters	Values	Targets / Source
Infection exit rate, π_X	0.78	Expected infection duration: 18 days
Unconstrained death rate,		Fatality rates in South Korea
$\quad \delta_1^u \times 100$	0.08	
$\quad \delta_2^u \times 100$	0.85	
$\quad \delta_3^u \times 100$	8.47	
Untreated death rate, δ_j^c	$2\delta_j^u$	Piguillem and Shi (2020)
Hospital capacity, κ	0.01	See discussion above
Transmission parameters,		Basic reproduction number, $R_0 = 2.2$,
$\quad$ consumption-related, β_c	0.23	and initial transmission equally
$\quad$ labor-related, β_ℓ	9.46	likely through three channels
$\quad$ other, β_e	0.57	
Flow value of life, $\bar{u}$	17.31	Value of statistical life: \$6.9 million (2019)
Disutility from not working, $\tilde{u}$	0.82	15 percent reduction in employment
		during pandemic
Efficiency units, η_{jI}	$0.5\eta_{jS}$	See discussion above

Sensitivity analysis in Appendix D reveals that the main results are robust to assuming, alternatively, a 30 percent loss in efficiency.

3.2 Calibration of the transition path

In this subsection, I start by discussing the initial conditions, the timing of a vaccine and cure, and transition path assumptions. I then discuss how the mitigation policies are set to align with mitigation policies implemented in the United States. Finally, I discuss how the time-varying transmissibility parameter ϵ_t is set.

Initial conditions and timeline. The economy starts in the pre-pandemic steady state in period $t = 0$. Then, in period $t = 1$ (March 27, 2020), the virus is introduced into the model so that 0.5 percent of the population is infected.[24] I assume that a vaccine becomes available in $t = 20$ (December 18, 2020), after which individuals over the age of 45 transit

---

[24]In the US, there were 17,982 COVID-19-related deaths during the 14-day period from March 27, 2020, to April 9, 2020, according to https://www.worldometers.info/coronavirus/country/us/. An initial infected rate of 0.5 percent generates deaths in the model ($t = 1$) that are consistent with the data.

to the recovered state with a probability of 10 percent every two weeks (5 percent for young individuals). This implies an expected wait time of 20 and 40 weeks for older and younger individuals, respectively. I assume that a cure is available in $t = \hat{t} \equiv 53$ (March 25, 2022), when all remaining susceptible and infected individuals transit to the recovered state with probability 1.[25]

While the steady state analysis is done in general equilibrium, the transition path analysis is done in partial equilibrium. In other words, factor prices (w, r), retirement income (s), and the labor income tax (τ_ℓ) are fixed at their pre-pandemic steady state levels.[26] Furthermore, the measure of newborns and their wealth distribution is also assumed to be constant throughout the transition. I provide a formal definition of the equilibrium with transition in Appendix B.

Mitigation policies. While counterfactual mitigation policies are the focus of the next sections, here they need to be specified to closely mimic mitigation policies implemented in the United States for calibration purposes. I focus on the set of policies implemented at the federal, state, and local levels that are relevant for virus mitigation. For example, I do not consider tax credit payments that may have been important for distributional consequences but otherwise not directly related to mitigating the virus.

The stay-at-home subsidy most closely resembles the $600 supplement to unemployment benefits (Federal Pandemic Unemployment Compensation, FPUC) that was in effect March 27, 2020, to July 27, 2020, and the $300 supplement in effect between December 26, 2020, and September 6, 2021 (Consolidated Appropriations Act and American Rescue Plan Act). While the stay-at-home subsidy has similarities to the FPUC unemployment supplement, one important difference is that unemployment benefits typically require involuntary unemployment, while the model subsidy is based on voluntary nonemployment. However, this difference is mitigated by the Pandemic Unemployment Act (PUA), which expanded the eligibility guidelines to include any individual out of work because of the pandemic, including

---

[25]This of course is counterfactual. However, the focus of this paper is optimal mitigation policies in the first stage of the pandemic—before vaccines became widely available in early 2021. There are many other papers that focus on the second stage of the pandemic (see, for example, Glover et al. 2022).

[26]While these assumptions have been made for computational tractability, there are also other considerations. For instance, Social Security benefits and contribution rates do not typically respond to recessions at a high frequency (and did not during the pandemic). One could also make the argument that wages did not change much during the first phase of the pandemic (before vaccines became widely available in early 2021) after accounting for composition effects (for example, see Rouse and Gimbel 2021).

Figure 1: Mitigation policies

(a) Subsidy

(b) Outside hours cap
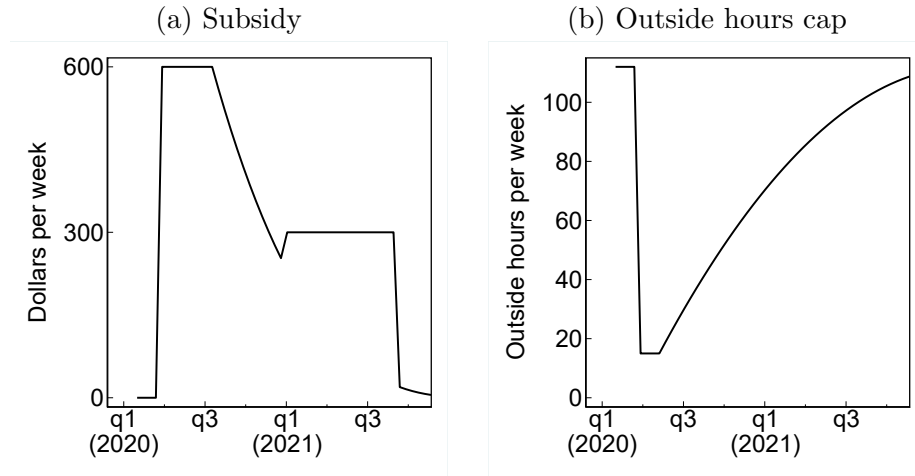
self-employed and gig workers, whose labor supply decisions are likely more voluntary.[27]

Based on the timeline of the FPUC, the model subsidy of \$600 per week begins at $t = 1$ (March 27, 2020), with a gradual reduction after $t = 9$ (July 30, 2020). There is a second subsidy of \$300, which begins at $t = 21$ (January 1, 2021) and ends at $t = 38$ (September 9, 2021). Figure 1 (panel a) plots the model subsidy amount over time. For the purpose of calibration, I set the consumption tax $\tau_c = 0$, since there was no associated tax increase in the data, implying that the government budget constraint does not hold. In contrast, in the counterfactual exercises in Section 5, I use the consumption tax as a means to fund the stay-at-home subsidy.

The model stay-at-home order (lockdown) most closely resembles the local- and state-level stay-at-home orders, which began in most states between March 23 and April 1. The duration and intensity varied substantially across locations and many areas began reopening in early May. One caveat is that the stay-at-home orders affected workers differently based on whether their place of work was essential (grocery stores) or social-intensive (restaurants and bars). By abstracting from sectors and occupations, the model cannot speak directly to these differences. However, to the extent that social-intensive occupations tend to have lower wages as documented by Kaplan et al. (2020), the model indirectly captures these differences since the lockdown disproportionately affects lower-wage individuals.

The model lockdown is implemented by restricting outside labor supply to less than $\bar{\ell}^o$ = 0.13, equivalent to 15 hours per week, beginning at $t = 1$ (March 27, 2020), and gradually

---

[27]To a lesser extent, the subsidy is also related to the Paycheck Protection Program (PPP), part of which was used by firms adversely impacted by the pandemic and mitigation policies to pay furloughed workers.

phasing out after $t = 4$ (May 8, 2020). In equilibrium, the cap is no longer binding for any individual after August 28, 2020. Figure 1 shows the time series of both the stay-at-home subsidy and the hours cap implemented in the model. Since the lockdowns implemented in the US differed across sectors and localities in scope and intensity, the model lockdown should roughly be interpreted as capturing an average effect.

Time-varying transmissibility. Finally, I discuss how the time-varying parameter that governs the virus transmission outside of consumption and work, ϵ_t, is set.

Because the progression of the pandemic in the data is a function of activities endogenous in the model and a host of other factors that are not in the model, I use the time-varying parameter to parsimoniously capture all of those other factors, which allows the model to generate exactly the time series of deaths observed in the United States. I use the first 12 months of data on biweekly deaths in the United States (March 27, 2020–March 25, 2021) to calibrate ϵ_t.[28] This corresponds to the first phase of the pandemic (before vaccines become widely available), which is the period of focus in this paper. I then set $\epsilon_t = 0$ for $t \geq 27$ (after March 26, 2021). Figure 2 plots the calibrated path of $\beta_e + \epsilon_t$, which governs the transmission of the virus other than through outside consumption and labor. It shows that the fitted ϵ_t is consistent with a) a seasonal variation in which the transmissibility declines during the summer months and b) changes in behavior not in the model such as increased transmission around holidays such as July 4, 2020.[29]

A final note about the transition path: While susceptible workers optimally prefer to work inside such that $\ell_i = \bar{\theta}_j(\varepsilon)\ell$, infected and recovered workers are indifferent about the composition of their hours. Thus, I assume that infected workers also prefer to work inside such that $\ell_i = \bar{\theta}_j(\varepsilon)\ell$, and recovered workers choose their hours to be entirely outside such that $\ell_i = 0$.

---

[28]See Appendix C.2 for details on the computation of the transition path, including how the path of ϵ_t is found.

[29]See, for example, Grassly and Fraser (2006) for a discussion of what causes seasonality in infectious diseases. One factor that may contribute to higher transmissibility in the winter is lower humidity, which increases the survival of the influenza virus in air. Another factor is seasonal changes in human behavior, such as those affected by summer vacation as well as a greater likelihood of family gatherings over the November and December holidays.

Figure 2: Time-varying transmissibility
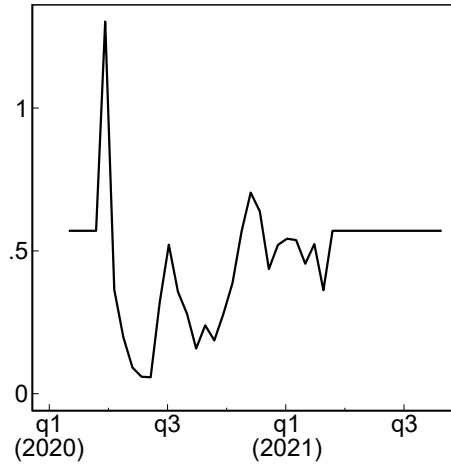
3.3 Model validity

In this subsection, I investigate the model's fit by examining both targeted and non-targeted moments in the steady state and during the course of the pandemic. The goal is to assess whether we can plausibly use the model as a laboratory to run counterfactual experiments.

Pre-pandemic steady state. Table 3 reports some steady state moments that illustrate the calibration's overall performance. In addition to successfully matching the targeted moments, the model also generates non-targeted moments that are reasonably close to the data. For example, the model generates a gini for disposable labor income that is very close to that in the data and a consumption gini that is also reasonably close. The model balances a wealth gini that is somewhat lower than in the data with a 75-to-25 ratio that is somewhat higher than in the data. There are two points to make regarding the wealth distribution. On the one hand, the model inherits the limitations in generating a sufficiently skewed wealth distribution that are common in standard incomplete market models. On the other hand, for the context of this paper, this limitation is mitigated by the fact that the behavior of wealthy individuals in this class of models is not very different from those of extremely wealthy individuals.

Transition path with pandemic. So far, we have confirmed that the model generates a reasonable starting point for the pandemic along the relevant margins. Another test of the model lies in its performance along the transition path when COVID-19 is introduced. Figure 3, which plots the time series of relevant variables for the model and data, shows that

Table 3: Targeted and non-targeted moments

	Data	Model
Targeted moments		
wealth/GDP	4.8	4.8
average weekly hours	34.4	34.4
average VSL (multiples of annual consumption per capita)	238.8	238.8
Non-targeted moments		
disposable earnings gini	0.37	0.36
consumption gini	0.33	0.25
wealth gini	0.74	0.59
wealth p75/p25	11.9	13.2

Notes: Statistics related to the disposable earnings, consumption, and wealth distribution in the data are computed on a per capita basis. See Appendix A.1 for details.

the model performs reasonably well in this dimension.

Panel (a) of Figure 3 shows that output contracts sharply in the second quarter with a strong recovery in the third quarter (as in the data). From there, the model deviates slightly from the data in generating small declines in output for two consecutive quarters before rebounding in the second quarter of 2021 whereas the recovery in the data is smoother. Panel (b) shows that consumption contracts less sharply relative to the data in the second quarter and rebounds less strongly in the third quarter.[30] Model consumption stalls in the fourth quarter and begins recovering in early 2021, as in the data. Overall, the model generates dynamics for aggregate output and consumption that are reasonably similar to the data.

Panels (c) and (d) of Figure 3 show the time series of outside consumption and labor. For each of these series, there is not a perfect data counterpart because of data limitations. Having said that, the model-generated series for outside consumption lies between the data series for Google mobility (retail and recreation) and OpenTable reservations. These are imperfect yet informative measures of consumption activities that are risky in the sense that they involve leaving the house and potentially getting the disease from or passing it onto

---

[30]One reason that model consumption declines less sharply than in the data may be that the model lockdown is implemented as a cap on outside hours, without a corresponding cap on outside consumption. In reality, business closures implied less opportunities for both outside labor and outside consumption.

Figure 3: Aggregates during the pandemic

(a) Output

(b) Consumption

(c) Outside consumption

(d) Outside hours

(e) Cases

(f) Deaths

(g) Share of deaths, 65+
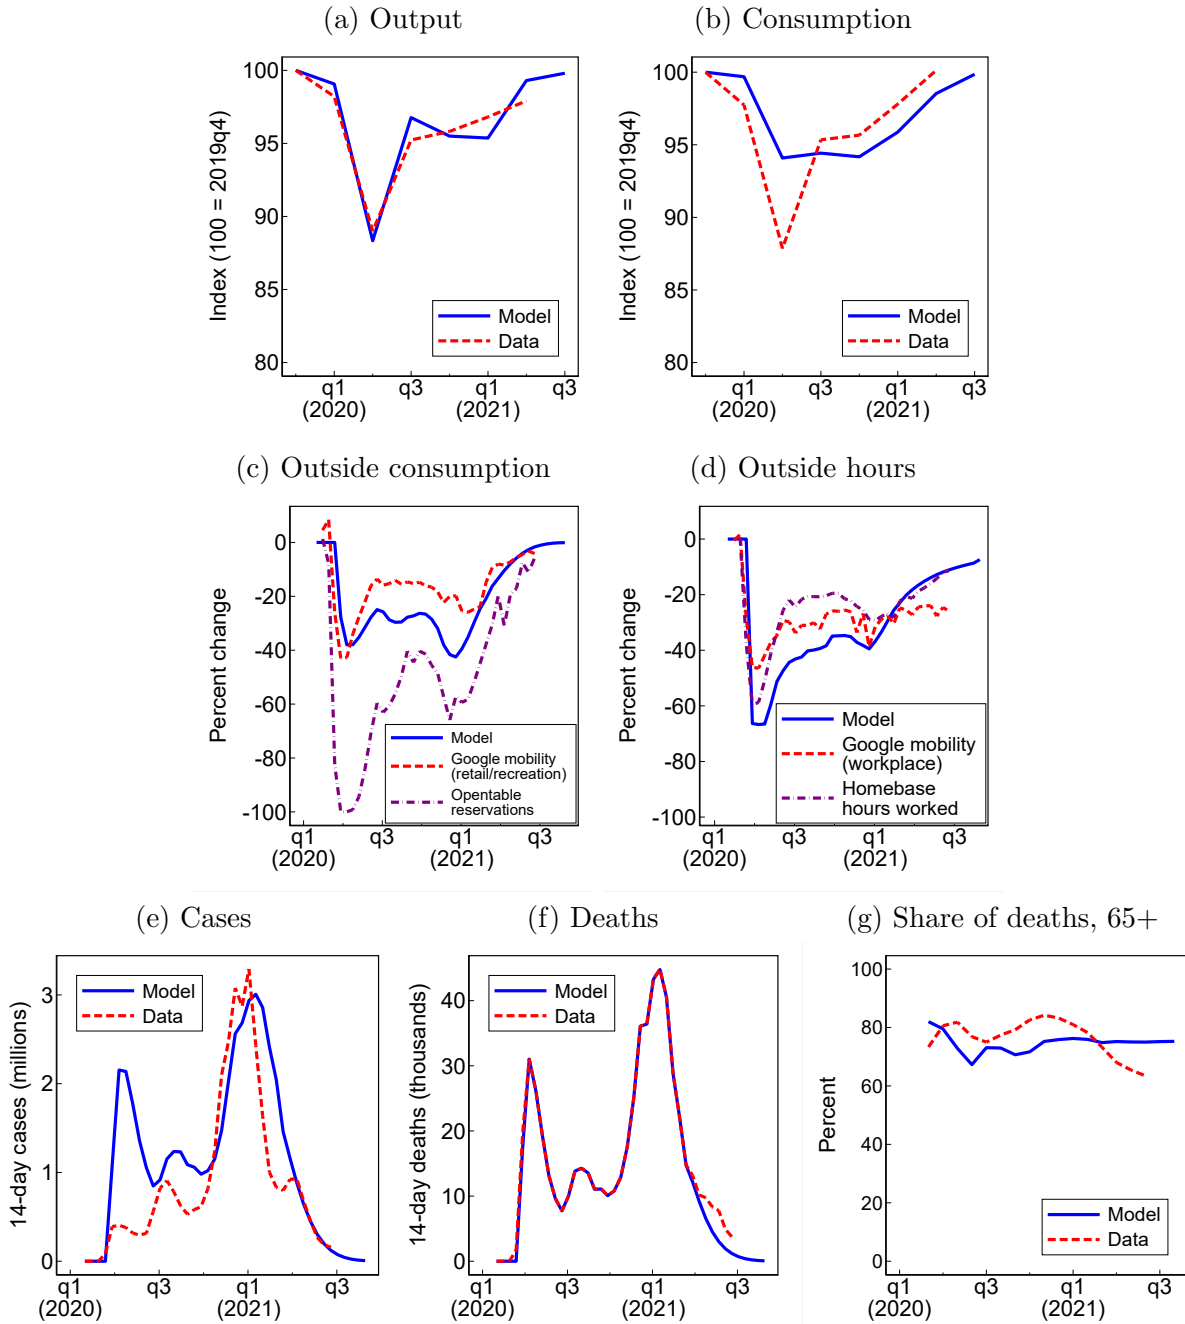
Notes: Both output and consumption in the data are linearly detrended at 2 percent per year. Outside consumption and hours in the model are relative to the pre-pandemic steady state. Google mobility and OpenTable reservations are year-over-year percent changes. Homebase hours are percent changes relative to the median for each day of the week during the January 4–31, 2020.

others. Similarly, the model generated series for outside hours tracks fairly closely the data series for Google mobility (workplace) and Homebase hours worked, which reports changes in hours worked for hourly employees that mostly work in contact-intensive sectors (e.g. restaurants, retail, health care, and other services). These measures are also imperfect but provide information regarding the changes in work-related activities that are risky because they take place outside of one's house. The model decline in outside hours is slightly larger than the data counterparts: Most of this can be explained by the fact that the pre-pandemic steady state assumes that no labor is done from home, whereas in the data, there was at least some labor done from home before the pandemic. Nonetheless, the dynamics for both outside consumption and labor are very similar to the data with large contractions in March 2020, followed by a strong recovery and a fourth-quarter decline that persisted into early 2021.

Panels (e) and (f) of Figure 3 plot the times series of cases and deaths, respectively. Recall that the time-varying transmissibility parameter ϵ_t was calibrated to fit the death series from March 27, 2020, to March 25, 2020. The data series for cases is well below the model generated series in the early part of the pandemic, possibly reflecting the widely reported lack of testing in early stages of the pandemic in the United States.[31] Finally, panel (g) of Figure 3 plots the share of COVID-19 deaths that are attributable to individuals whose ages are 65 or older. The model captures the fact that the vast majority of deaths from COVID-19 are among older individuals.

Overall, the model generates times series for relevant variables that are reasonably similar to the data. Having established that the model can plausibly be used as a laboratory for running counterfactual experiments, I quantify the effects of the pandemic and the heterogeneous mitigation efforts across the distribution and examine optimal policies in the next two sections.

4 Quantitative exercises

This section uses the calibrated model to investigate the properties of private mitigation and the distributional consequences of the pandemic and mitigation policies. First, I study how private mitigation differs across the age, income, and wealth distribution. Second, I quantify the effects of the pandemic and the mitigation policies that resemble those implemented in

---

[31]Note that the model projects a monotonic decline in cases and deaths from January 2021 onward. This, of course, is counterfactual. Recall, though, that this paper focuses on the first phase of the pandemic, before vaccines became widely available.

the US. In particular, I contrast the *stay-at-home order* (lockdown) and the *stay-at-home subsidy*. While both policies reduce infections and deaths, the subsidy policy delivers a higher welfare gain and is favored by all individuals in the economy, whereas the lockdown benefits older individuals at the expense of low-wage workers.

4.1   Private mitigation

To study how private mitigation differs across the age, income, and wealth distribution, consider the plots of outside consumption and labor—in the absence of any public mitigation policies (Figure 4). The decline in outside consumption and hours is broad based. However, the decline in outside consumption is much greater for middle-aged and old agents than for young agents (panels a–c), and the decline in outside hours is much larger for middle-aged workers than for young workers (panels d–e). As a result, young workers experience a much larger increase in infections, as shown in panel (f).[32]
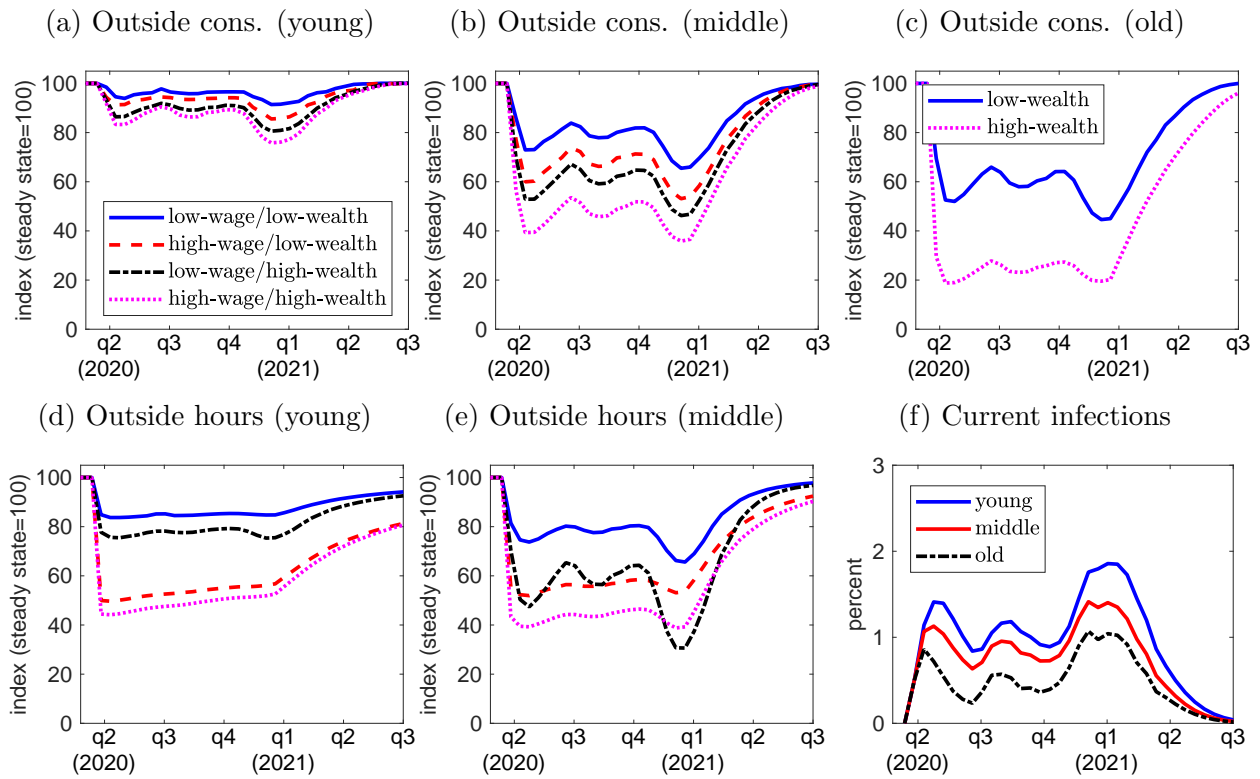
Additionally, within each age cohort, the declines in outside consumption and hours are the smallest for low-wage/low-wealth workers. The reasons are threefold. First, while high-wage workers can switch to mostly working from home, low-wage workers are unable to do so. Second, low-wealth workers are not able to weather prolonged time away from work. The fact that low-wage/high-wealth workers significantly reduce their outside hours during the pandemic (even in the absence of mitigation policies) suggests that the lack of precautionary savings prevents low-wage/low-wealth workers from similarly reducing their labor supply. Finally, high-wage and high-wealth individuals are more willing to curtail activities that increase infection risk because they have higher expected continuation values that are derived from higher expected future consumption and leisure.[33] Thus, they are more willing to substitute outside consumption with inside consumption, increasingly so during the peaks of the pandemic.

These properties are qualitatively consistent with the data. Using the Real-Time Population Survey (RPS), an individual-level nationwide survey developed by Bick et al. (2020), I document that, from February 2020 to May 2020, low-income individuals reduced their outside labor (days commuted) by 36 percent, which is significantly less than the 47 percent reduction among their high-income counterparts. These properties are also qualitatively

---

[32]Despite lower infection rates, because older individuals face a higher fatality risk, they account for the vast majority of all COVID deaths, as in the data.

[33]For example, the model-implied VSL for high-wage/high-wealth individuals is roughly six times higher than that for low-wage/low-wealth individuals.

Figure 4: Response to pandemic (no public mitigation)

(a) Outside cons. (young)

(b) Outside cons. (middle)

(c) Outside cons. (old)

(d) Outside hours (young)

(e) Outside hours (middle)

(f) Current infections

Notes: Low wage and high wage correspond to below and above the steady state median wage, respectively.
Low wealth and high wealth correspond to below and above the steady state median wealth, respectively.

consistent with empirical facts documented by recent papers. Brotherhood et al. (2020a) use Brazilian data to find that social distancing increased by less in areas with slums (low-income) than those without slums (high-income). Chetty et al. (2020) find that spending and mobility declined more in higher-income locations than in lower-income locations. Hevia et al. (2022) use data from Bogotá to find that individuals in the low socioeconomic stratum (SES) are twice as likely to have been infected with COVID-19 than high SES individuals. Similarly, Bruckner et al. (2021) use a study of adults in Orange County, California, to find higher COVID-19 prevalence among low-income adults relative to high-income adults and among young adults relative to old adults.

4.2   Welfare consequences of COVID-19 and mitigation policies

The previous subsection highlighted the externalities at work: Young workers—especially those that are low-wage and low-wealth—do not reduce their outside consumption and labor as much as their older counterparts and incur higher rates of infection. These responses are individually rational in the sense that young workers do not face high fatality risk. However, higher rates of infection among young agents contribute to higher infections among older individuals, who face higher fatality rates. In this subsection, I quantify the distributional effects of the pandemic and mitigation policies.

Table 4 shows the welfare, economic, and health outcomes of the pandemic and mitigation policies, decomposed by the effects of the subsidy policy and the lockdown policy. The first column shows that even without mitigation policies, the contraction of output is large, amounting to a nearly 3 percent loss in two-year output. The welfare consequences of the pandemic—measured as permanent consumption equivalents—are large across the distribution, but largest for retirees, who face a higher fatality risk.

Mitigation policies that most closely resemble mitigation policies implemented in the US (described in Section 3.2) reduce deaths and output, relative to no public mitigation (second column of Table 4). The third and fourth columns of Table 4 decompose the effects of the subsidy and lockdown separately. Both policies improve average welfare and reduce deaths. All individuals benefit from the subsidy policy (i.e. the policy is a Pareto improvement). This is not surprising given that the benefits of the subsidy are provided without a corresponding cost (i.e. no tax increase). In the next section, however, I show that Pareto improving policies are possible even when the subsidies are fully financed with a consumption tax. In contrast, the lockdown policy is favored by only 72.7 percent of the initial population. This is because the lockdown is mainly favored by older agents who most value the lower risk of death

Table 4: Welfare consequences of pandemic and mitigation policies

	no public mitigation	US mitigation	subsidy only	lockdown only
welfare	−5.4	−4.4	−4.4	−5.3
working-age	−3.1	−2.4	−2.4	−3.1
retired	−14.8	−12.4	−12.6	−14.2
low-wage	−1.8	−1.3	−1.3	−2.0
high-wage	−4.3	−3.5	−3.6	−4.1
low-wealth	−3.8	−3.0	−3.0	−3.9
high-wealth	−6.9	−5.7	−5.8	−6.6
policy support		89.7	100.0	72.7
2-year output	97.3	96.7	97.7	96.4
deaths per 10k	22.0	17.7	18.0	21.0

Notes: Low- and high-wage correspond to below and above the median wage, respectively. Low- and high-wealth corresponds to below and above the median wealth, respectively. Welfare refers to permanent consumption equivalents (i.e. the percent change in permanent consumption in the pre-pandemic steady state that would make an individual indifferent with the transition path with the pandemic from $t = 1$ forward). Blue and red colors denote groups with welfare gains and losses from the mitigation policies, respectively. Policy support refers to the percent of the initial population that benefits from the mitigation policy. Output refers to output from $t = 1$ (March 27, 2020) to $t = 52$ (March 24, 2022), compared with the analogous 52-period output in the steady state, indexed at 100.

induced by the policy and is opposed by low-wage workers for whom the lockdown policy is most binding. In terms of output, the subsidy alone increases output relative to no public mitigation, whereas the lockdown is contractionary. In the next section, which examines the properties of optimal policies, I show that lockdowns generally are a very inefficient means to save lives and discuss how mitigation policies can simultaneously increase output and save lives.

5   Optimal mitigation policies

I investigate the properties of optimal mitigation policies over a limited set of policy instruments. In particular, I solve for the transition paths for 3,800 combinations of policy parameters, which vary along the weekly subsidy amount, T ($0–$1,200), the duration (4–14 months), and the speed at which the subsidy phases out[34], with and without lockdowns of varying intensities (10, 15, and 20 weekly hour limits on outside labor). In all cases with a positive subsidy, I solve for a consumption tax—levied over the same period as the subsidy—that clears the government budget constraint in present value, making all configurations budget-neutral. All other parameters, including the time-varying transmissibility, ϵ_t, are kept the same as described in Section 3.

Figure 5 plots the effects of varying subsidy amounts and duration, keeping fixed the speed at which the subsidy phases out and no lockdown. Panel (a) shows that average welfare is increasing in both the subsidy amount and duration and panel (b) shows that for subsidies that are not too large or long, support for the mitigation policy is unanimous. Falling support for the policy is due to the higher consumption tax rate associated with larger and longer duration subsidies (panel e). Panel (c) demonstrates that deaths are decreasing in both the subsidy amount and duration.
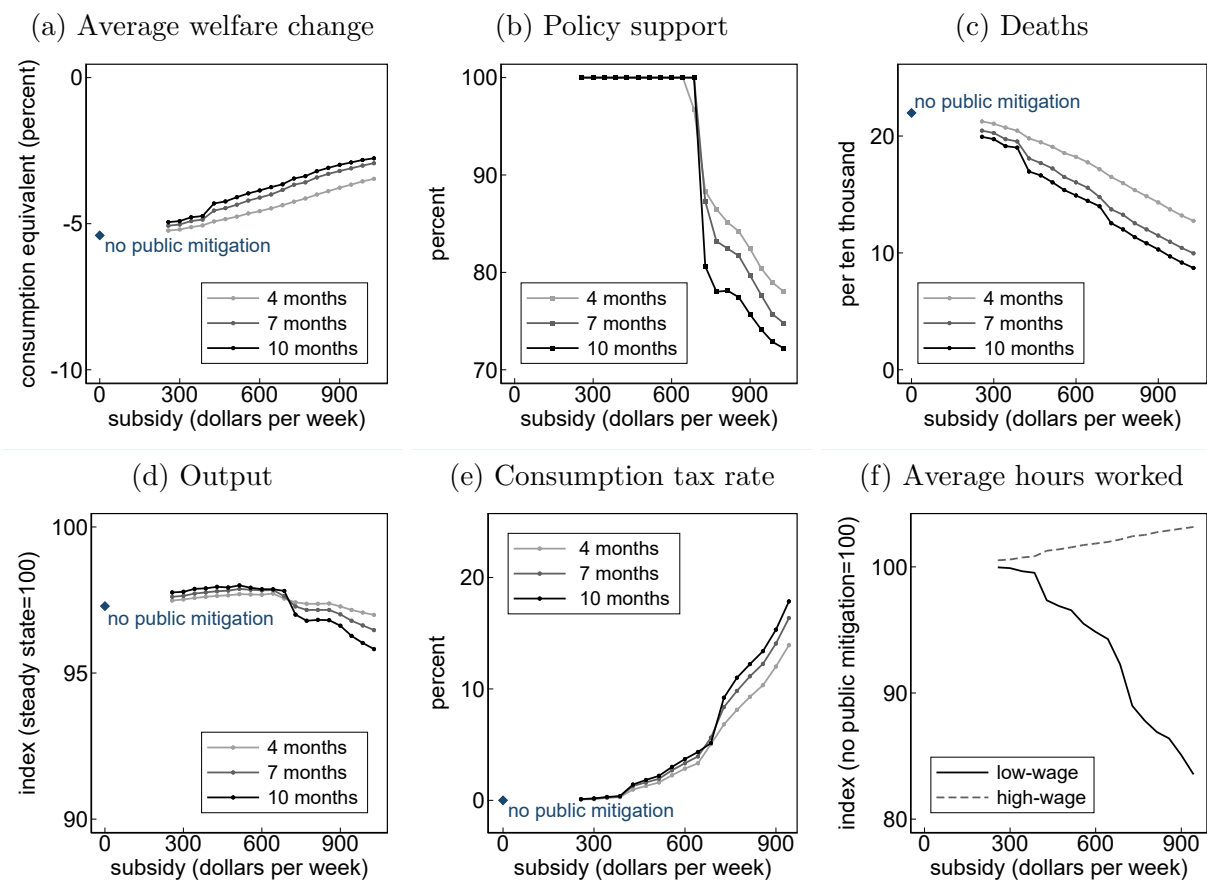
Interestingly, the effects of the subsidy amount on output are nonmonotonic (panel d). In particular, for moderate subsidy amounts, the mitigation policy actually increases output

---

[34]Specifically, the subsidy amount follows

$$
T_t = \begin{cases} T & \text{if } 1 \leq t \leq \bar{t} \\ T\left(1 - \dfrac{t - \bar{t}}{\hat{t} - \bar{t}}\right)^x & \text{if } \bar{t} < t < \hat{t} \\ 0 & \text{otherwise.} \end{cases} \tag{19}
$$

where $\bar{t}$ denotes the period at which the phase-out begins and $x \in \{1, 2, 4, 8\}$ determines the speed of the phase-out. The consumption tax and lockdowns, where applicable, are phased out in an analogous manner.

Figure 5: Effects of subsidy amount and duration

(a) Average welfare change

(b) Policy support

(c) Deaths

(d) Output

(e) Consumption tax rate

(f) Average hours worked

Notes: The graphs show the effects of varying subsidy amounts and duration, keeping fixed the pace at which the subsidy phases out and no lockdown. Average welfare change reports the population-weighted average of individual consumption equivalents. Policy support refers to the percent of the initial population that benefits from the mitigation policy. Output refers to output from $t = 1$ (March 27, 2020) to $t = 52$ (March 24, 2022), compared with the analogous 52-period output in the steady state, indexed at 100. Average hours worked refers to hours from $t = 1$ (March 27, 2020) to $t = 52$ (March 24, 2022), compared with the analogous 52-period average hours with no public mitigation, indexed at 100, for low- and high-wage individuals (i.e., individuals with below and above the median wage, respectively).

relative to the case with no public mitigation. This is due to two opposing effects. The first is the *direct effect* of the subsidy reducing labor supply, holding fixed the severity of the pandemic. Moderate subsidies (less than $500 per week) induce only the very low-wage workers to reduce their hours, leading to a small decline in output. The direct effect can be seen in panel (f), where larger subsidies induce larger reductions in hours of low-wage workers, relative to no public mitigation. The second is the *indirect effect* of the subsidy attenuating the pandemic, thereby leading to increased labor supply. Specifically, because the labor from low-wage workers is almost exclusively "outside," most of the reduction in hours due to the direct effect meaningfully contributes to mitigating the virus, making it safer to engage in more economic activities.[35] This indirect effect is most evident among high-wage workers, as can be seen in panel (f), which demonstrates that larger subsidies induce more hours of high-wage workers, relative to no public mitigation. At moderate subsidy amounts, the indirect effect dominates the direct effect, leading to an increase in output relative to the no-public-mitigation scenario.

Mitigation policies have often been portrayed in the context of a tradeoff between output and health. Figure 6 illustrates the relation between two-year output and lives saved: The solid line depicts the best mix of policies that maximize output for each level of lives saved during the pandemic and the dashed line is the analogous line that does so with only lockdowns. We can see that, when exclusively relying on lockdowns to save lives, there is indeed a tradeoff between output and lives saved: A better health outcome is associated with lower output. However, we can also see that there are many policy configurations that simultaneously increase output and save more lives, relative to no public mitigation. In fact, Figure 6 shows that COVID-19 deaths can be reduced by as much as 45 percent without any further reduction in output, relative to no public mitigation. As reference points, I include the best Pareto-improving policy, which maximizes welfare conditional on full support; the output-maximizing configuration; and the configuration that most closely resembles US mitigation policy as described in Section 3.2. The policy parameters of the best Pareto-improving and output-maximizing configurations are summarized in Table 5 (top panel).

Relative to no public mitigation, Table 5 shows that a subsidy of $700 per week for 10 months can lead to a substantial reduction in deaths and increased output. In addition, a slightly smaller subsidy of $500 per week for 10 months increases output even further (but saves fewer lives). Neither the best Pareto-improving nor the output-maximizing policy

---

[35]Additionally, taxes used to fund the subsidy also contribute to mitigating the virus, allowing additional economic activity.

Figure 6: Output and lives
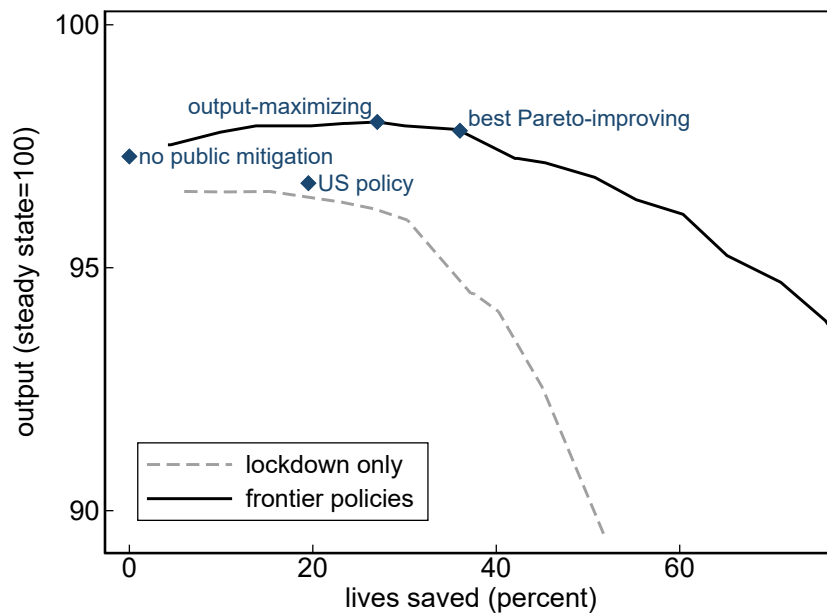
Notes: The solid line depicts the best mix of policies that maximize output for each level of lives saved during the pandemic and the dashed line is the analogous line that does so with only lockdowns. Output refers to output from $t = 1$ (March 27, 2020) to $t = 52$ (March 24, 2022), compared with the analogous 52-period output in the steady state. Lives saved are relative to no public mitigation.

Table 5: Optimal mitigation policies

	no public mitigation	output-maximizing policy	best Pareto-improving	subsidy only	tax only
subsidy ($/week)	0	500	700	700	0
duration (months)	0	10	10	10	10
cons. tax (percent)	0	2.2	5.5	0	5.5
lockdown	no	no	no	no	no
welfare	−5.4	−4.1	−3.7	−3.7	−5.4
working-age	−3.1	−2.2	−1.9	−1.9	−3.1
retired	−14.8	−11.5	−10.5	−10.8	−14.5
low-wage	−1.8	−1.2	−0.9	−0.8	−1.9
high-wage	−4.3	−3.3	−3.0	−3.0	−4.3
low-wealth	−3.8	−2.8	−2.4	−2.3	−3.9
high-wealth	−6.9	−5.4	−4.9	−5.0	−6.8
policy support		100.0	100.0	100.0	43.4
2-year output	97.3	98.0	97.8	97.8	97.4
deaths per 10k	22.0	16.0	14.0	15.1	20.9

Notes: Average welfare change reports the population-weighted average of individual consumption equivalents. Output refers to output from $t = 1$ (March 27, 2020) to $t = 52$ (March 24, 2022), relative to the analogous 52-period output in the steady state, indexed at 100.

features a lockdown. The middle panel of Table 5 shows the welfare consequences of the mitigation policies across the distribution. In particular, the second and third columns reveal that the welfare gains from both the output-maximizing and the best Pareto-improving policies (derived by subtracting the first column from the second or third column) are large and widely shared. Importantly, note that both of these policies are Pareto improvements.

The last two columns of Table 5 decompose the effects of the best Pareto-improving subsidy and tax separately (I omit the analogous decomposition for the output-maximizing policy as the effects are smaller but qualitatively similar). The majority of the lives saved and the welfare gains from the optimal policy are due to the subsidy. However, the consumption tax also contributes to saving lives and increasing output (by making it safer to work); it also results in a welfare gain for older and wealthy individuals who put a larger value in

containing the virus. In Appendix D, I show that using a smarter Pigouvian tax on outside consumption is even more effective in savings lives and improving welfare.[36]

Discussion and caveats. The stay-at-home subsidy, fully funded by a consumption tax, can be Pareto improving because it addresses a key externality: Susceptible individuals understand that outside labor and consumption increases their own risk of infection, but they do not take into account the risk induced by the outside labor and consumption of their future infected selves. As seen in Section 4.1, private mitigation is smallest for young, low-wage, and low-wealth individuals. This subsidy is particularly effective because it provides the incentive for this demographic to stay at home, attenuating the externality and consequently the severity of the pandemic.

Might a similar tax-and-subsidy policy be Pareto improving in normal times? When conducting the analogous policy search over the same set of instruments in the absence of a pandemic, I am unable to find any policy that is welfare improving, let alone one that is Pareto improving (see Appendix D). This suggests that the externalities generated by the COVID-19 pandemic are crucial for the possibility of allowing Pareto improving policies.

In sensitivity analysis (Appendix D), I show that the main results of the paper are robust to alternative values for $\bar{u}$, which governs the value of statistical life, η_{jI}, which governs the efficiency loss during infection, the disutility of not working ($\tilde{u}$), the basic reproduction number (R_0), and the preference for inside consumption (γ). The results are also robust to many alternative model assumptions. For example, in Appendix D, I show that the main results are robust to eliminating the hospital capacity constraint, funding the subsidy using a tax on outside consumption, funding the subsidy using a tax on labor income, and an alternative bequest assumption where a deceased individual is replaced by a newborn with the same productivity and wealth level.

The results could be sensitive, however, to other modeling assumptions. For instance, the model assumes that both outside and inside consumption use an identical good, differentiated only in the manner in which that good is consumed (e.g. food at a restaurant versus food at home). If, alternatively, inside and outside consumption were different goods (with limited labor mobility across sectors), then a reduction in the relative demand for outside consumption would lead to declines in the relative price of outside consumption

---

[36]I do not consider this for the baseline, however, as implementing such a Pigouvian tax could be difficult, since the relevant margin is not the physical location of consumption but whether the activity is high-risk for transmission. For example, in-home dining can be high-risk if it includes individuals from outside of one's own household.

and the relative wage in that sector. There are two points to make here. First, including the dynamics of the relative price of outside consumption goods would likely strengthen the need for mitigation policies, since private mitigation would be reduced as a result of cheaper outside consumption. Second, while including labor immobility across sectors and the dynamics of the relative wage would certainly change the distribution of welfare losses from the pandemic, it would not necessarily change the importance of the subsidy-and-tax mitigation policy since the low-wage and low-wealth individuals that accept the subsidy are still better off relative to no public mitigation (by revealed preference), and higher-wage and higher-wealth individuals would still value the attenuated severity of the pandemic.

The model also assumes that inside labor and outside labor are perfect substitutes as factor inputs. What if, alternatively, the production of outside consumption goods was comparatively intensive in outside labor (e.g. restaurant dining requires onsite restaurant staff)? One concern could be that subsidizing low-wage individuals (who work almost exclusively outside) to stay at home would create labor shortages in the production of outside goods. All else equal, this would lead to an increase in the relative cost of outside labor and the relative price of outside goods, the latter of which would be desirable from a policy perspective since it would contribute to reducing the transmission of the virus (it would be akin to a Pigouvian tax on outside consumption). Thus, incorporating limited factor substitutability would also likely strengthen the need for mitigation policies.

A major caveat to the main findings of this paper is that the model abstracts from children. As shown by Agostinelli et al. (2022) and Fuchs-Schündeln et al. (2020), the COVID-19 pandemic and associated school closures have a first order welfare effect on school-age children. While children do not directly contribute to economic production, they can play an important role in transmitting the virus to others. Thus, an interesting avenue for future research would be to jointly study optimal mitigation and school closure policies.

Finally, the model omits other potentially important sources of heterogeneity, such as pre-existing health conditions and finer age bins. For example, as studied by Bairoliya and Imrohoroglu (2020), lockdowns that specifically target age and pre-existing health status can be an effective tool in combating the virus. Thus smarter and better targeted lockdowns may affect the size of the optimal subsidy-and-tax policy. Investigating optimal mitigation policies that incorporate these and other additional sources of heterogeneity are left for future research.

6 Conclusion

In this paper, I developed a quantitative life cycle–economic epidemiology model that was used to measure the heterogeneous welfare consequences of COVID-19 and to investigate the properties of optimal mitigation policies. Using the calibrated model, I show that private mitigation plays a substantial role in reducing deaths and output during the pandemic, but there is large heterogeneity in the intensity of private mitigation across individuals. In particular, reductions in economic activities that contribute to the spread of COVID-19 are larger for individuals that are older, have higher wages, or have higher wealth. I also show that stay-at-home subsidies are superior to stay-at-home orders because they reduce deaths by more and output by less. Furthermore, Pareto improving policies can simultaneously improve economic and public health outcomes.

Widespread vaccines have contributed to putting the pandemic in the rear-view mirror in the United States and other developed economies, though the advent of variants that bypass existing vaccines remains a concern. Thus, the policy implications raised in this paper are valuable not only for pandemics in the future, but also in the present, especially for other countries where wide-scale vaccine implementations may take longer.

References

ACEMOGLU, D., V. CHERNOZHUKOVZ, I. WERNING, AND M. D. WHINSTON (2020): "Optimal Targeted Lockdowns in a Multi-Group SIR Model," Working Papers 826, Economic Research Southern Africa, https://ideas.repec.org/p/rza/wpaper/826.html.

AGOSTINELLI, F., M. DOEPKE, G. SORRENTI, AND F. ZILIBOTTI (2022): "When the great equalizer shuts down: Schools, peers, and parents in pandemic times," *Journal of public economics*, 206, 104574, https://doi.org/10.1016/j.jpubeco.2021.104574.

ALLEN, F. AND D. GALE (1994): "Limited market participation and volatility of asset prices," *The American Economic Review*, 933–955, https://www.jstor.org/stable/2118039.

ALVAREZ, F., D. ARGENTE, AND F. LIPPI (2020): "A Simple Planning Problem for COVID-19 Lockdown," *Covid Economics*, 14, 1–32, https://doi.org/10.3386/w26981.

ARGENTE, D. O., C.-T. HSIEH, AND M. LEE (2020): "The Cost of Privacy: Welfare Effect of the Disclosure of COVID-19 Cases," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27220.

ASPRI, A., E. BERETTA, A. GANDOLFI, AND E. WASMER (2021): "Mortality containment vs. economics opening: optimal policies in a SEIARD model," *Journal of mathematical economics*, 93, 102490, https://doi.org/10.1016/j.jmateco.2021.102490.

ASSENZA, T., F. COLLARD, M. DUPAIGNE, P. FÈVE, C. HELLWIG, S. KANKANAMGE, AND N. WERQUIN (2020): "The Hammer and the Dance: Equilibrium and Optimal Policy during a Pandemic Crisis," CEPR Discussion Papers 14731, C.E.P.R. Discussion Papers, https://ideas.repec.org/p/cpr/ceprdp/14731.html.

ATKESON, A. (2020): "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w26867.

AUM, S., S. Y. T. LEE, AND Y. SHIN (2020): "Inequality of fear and self-quarantine: Is there a trade-off between GDP and public health?" *Covid Economics*, 14, 143–174, https://doi.org/10.3386/w27100.

AZZIMONTI, M., A. FOGLI, F. PERRI, AND M. PONDER (2020): "Pandemic Control in ECON-EPI Networks," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27741.

BAIROLIYA, N. AND A. IMROHOROGLU (2020): "Macroeconomic Consequences of Stay-At-Home Policies During the COVID-19 Pandemic," *Covid Economics*, 13, 71–90, https://www.marshall.usc.edu/sites/default/files/bairoliy/intellcont/macroeffects_of_corona-1.pdf.

BAQAEE, D., E. FARHI, M. J. MINA, AND J. H. STOCK (2020): "Reopening Scenarios," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27244.

BERGER, D. W., K. F. HERKENHOFF, AND S. MONGEY (2020): "An SEIR infectious disease model with testing and conditional quarantine," *Covid Economics*, 13, 1–30, https://doi.org/10.3386/w26901.

BICK, A., A. BLANDIN, AND K. MERTENS (2020): "Work from home after the COVID-19 Outbreak," Tech. rep., https://doi.org/10.24149/wp2017r1.

BIRINCI, S., F. KARAHAN, Y. MERCAN, AND K. SEE (2020): "Labor Market Policies during an Epidemic," Tech. rep., https://doi.org/10.20955/wp.2020.024.

BODENSTEIN, M., G. CORSETTI, AND L. GUERRIERI (2020): "Social distancing and supply disruptions in a pandemic," *Covid Economics*, 19, 1–52, https://ssrn.com/abstract=3594260.

BOGNANNI, M., D. HANLEY, D. KOLLINER, AND K. MITMAN (2020): "Economic Activity and COVID-19 Transmission: Evidence from an Estimated Economic-Epidemiological Model," Tech. rep., University of Pittsburgh, https://doughanley.com/files/papers/COVID.pdf.

BOPPART, T., K. HARMENBERG, J. HASSLER, P. KRUSELL, AND J. OLSSON (2020): "Integrated epi-econ assessment," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w28282.

BRODEUR, A., D. M. GRAY, A. ISLAM, AND S. BHUIYAN (2020): "A Literature Review of the Economics of COVID-19," Tech. rep., https://ssrn.com/abstract=3636640.

BROTHERHOOD, L., T. CAVALCANTI, D. DA MATA, AND C. SANTOS (2020a): "Slums and pandemics," Tech. rep., http://dx.doi.org/10.2139/ssrn.3665695.

BROTHERHOOD, L., P. KIRCHER, C. SANTOS, M. TERTILT, ET AL. (2020b): "An economic model of the Covid-19 pandemic with young and old agents: Behavior, testing and policies," Tech. rep., https://econpapers.repec.org/paper/ptuwpaper/w202014.htm.

Bruckner, T. A., D. M. Parker, S. M. Bartell, V. M. Vieira, S. Khan, A. Noymer, E. Drum, B. Albala, M. Zahn, and B. Boden-Albala (2021): "Estimated seroprevalence of SARS-CoV-2 antibodies among adults in Orange County, California," *Scientific reports*, 11, 1–9, https://doi.org/10.1038/s41598-021-82662-x.

Carroll, D. and S. Hur (2020a): "On the distributional effects of international tariffs," Tech. rep., Federal Reserve Bank of Cleveland, https://doi.org/10.26509/frbc-wp-202018.

Carroll, D. R. and S. Hur (2020b): "On the heterogeneous welfare gains and losses from trade," *Journal of Monetary Economics*, 109, 1–16, https://doi.org/10.1016/j.jmoneco.2019.10.009.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2007): "Business cycle accounting," *Econometrica*, 75, 781–836, https://doi.org/10.1111/j.1468-0262.2007.00768.x.

Chari, V. V., R. Kirpalani, and C. Phelan (2020): "The hammer and the scalpel: On the economics of indiscriminate versus targeted isolation policies during pandemics," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27232.

Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, et al. (2020): "How did covid-19 and stabilization policies affect spending and employment? a new real-time economic tracker based on private sector data," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27232.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011): "Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins," *American Economic Review*, 101, 471–75, https://doi.org/10.1257/aer.101.3.471.

Dingel, J. I. and B. Neiman (2020): "How many jobs can be done at home?" *Covid Economics*, 1, 16–24, https://doi.org/10.3386/w26948.

Dotsey, M., W. Li, and F. Yang (2014): "Consumption and time use over the life cycle," *International Economic Review*, 55, 665–692, https://doi.org/10.1111/iere.12066.

Duernecker, G. and B. Herrendorf (2018): "On the allocation of time–A quantitative analysis of the roles of taxes and productivities," *European Economic Review*, 102, 169–187, https://doi.org/10.1016/j.euroecorev.2017.10.025.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020): "The macroeconomics of epidemics," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w26882.

Fajgelbaum, P. D., A. Khandelwal, W. Kim, C. Mantovani, and E. Schaal (2021): "Optimal lockdown in a commuting network," *American Economic Review: Insights*, 3, 503–22, https://doi.org/10.1257/aeri.20200401.

Farboodi, M., G. Jarosch, and R. Shimer (2020): "Internal and external effects of social distancing in a pandemic," *Covid Economics*, 9, 22–58, https://doi.org/10.3386/w27059.

Fauci, A. S., H. C. Lane, and R. R. Redfield (2020): "Covid-19-navigating the uncharted," *New England Journal of Medicine*, 382, 1268–1269, https://doi.org/10.1056/NEJMe2002387.

Ferguson, N. M., D. A. Cummings, C. Fraser, J. C. Cajka, P. C. Cooley, and D. S. Burke (2006): "Strategies for mitigating an influenza pandemic," *Nature*, 442, 448–452, https://doi.org/10.1038/nature04795.

Fernández-Villaverde, J. and C. I. Jones (2020): "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities," Tech. rep., National Bureau of Economic Research, http://doi.org/10.3386/w27128.

Floden, M. and J. Lindé (2001): "Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?" *Review of Economic dynamics*, 4, 406–437, https://doi.org/10.1006/redy.2000.0121.

French, E. and J. B. Jones (2011): "The effects of health insurance and self-insurance on retirement behavior," *Econometrica*, 79, 693–732, https://doi.org/10.3982/ECTA7560.

Fuchs-Schündeln, N., D. Krueger, A. Ludwig, and I. Popova (2020): "The long-term distributional and welfare effects of Covid-19 school closures," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27773.

Garibaldi, P., E. R. Moen, and C. A. Pissarides (2020): "Modelling contacts and transitions in the SIR epidemics model," *Covid Economics*, 5, 1–20, https://www.carloalberto.org/wp-content/uploads/2020/04/garibaldi.pdf.

GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RIOS-RULL (2020): "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic," *Covid Economics*, 6, 22–64, https://doi.org/10.1006/redy.2000.0121.

GLOVER, A., J. HEATHCOTE, D. KRUEGER, J.-V. RIOS-RULL, ET AL. (2022): "Optimal age-based vaccination and economic mitigation policies for the second phase of the covid-19 pandemic," *Journal of Economic Dynamics and Control*, https://doi.org/10.18651/RWP2021-15.

GRASSLY, N. C. AND C. FRASER (2006): "Seasonal infectious disease epidemiology," *Proceedings of the Royal Society B: Biological Sciences*, 273, 2541–2550, https://doi.org/10.1098/rspb.2006.3604.

GREENWOOD, J., P. KIRCHER, C. SANTOS, AND M. TERTILT (2019): "An equilibrium model of the African HIV/AIDS epidemic," *Econometrica*, 87, 1081–1113, https://doi.org/10.3982/ECTA11530.

HALL, R. E., C. I. JONES, AND P. J. KLENOW (2020): "Trading Off Consumption and COVID-19 Deaths," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27340.

HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2014): "Consumption and labor supply with partial insurance: An analytical framework," *American Economic Review*, 104, 2075–2126, https://doi.org/10.1257/aer.104.7.2075.

HEVIA, C., M. MACERA, AND P. A. NEUMEYER (2022): "Covid-19 in unequal societies," *Journal of Economic Dynamics and Control*, forthcoming, https://doi.org/10.1016/j.jedc.2022.104328.

HUR, S. (2018): "The lost generation of the Great Recession," *Review of Economic Dynamics*, 30, 179–202, https://doi.org/10.1016/j.red.2018.05.004.

HUR, S. AND M. JENUWINE (2020): "Lessons on the Economics of Pandemics from Recent Research," *Economic Commentary*, https://doi.org/10.26509/frbc-ec-202011.

JONES, C., T. PHILIPPON, AND V. VENKATESWARAN (2020): "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home," Tech. rep., https://doi.org/10.3386/w26984.

KAPICKA, M. AND P. RUPERT (2020): "Labor markets during pandemics," Tech. rep., UC Santa Barbara.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary policy according to HANK," *American Economic Review*, 108, 697–743, https://doi.org/10.1257/aer.20160042.

———— (2020): "The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the US," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w27794.

KEHOE, T. J., K. J. RUHL, AND J. B. STEINBERG (2018): "Global imbalances and structural change in the United States," *Journal of Political Economy*, 126, 761–796, https://doi.org/10.1086/696279.

KERMACK, W. O. AND A. G. MCKENDRICK (1927): "A contribution to the mathematical theory of epidemics," *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115, 700–721, https://doi.org/10.1098/rspa.1927.0118.

KOPECKY, K. A. AND R. M. SUEN (2010): "Finite state Markov-chain approximations to highly persistent processes," *Review of Economic Dynamics*, 13, 701–714, https://doi.org/10.1016/j.red.2010.02.002.

KREMER, M. (1996): "Integrating behavioral choice into epidemiological models of AIDS," *The Quarterly Journal of Economics*, 111, 549–573, https://doi.org/10.2307/2946687.

KRUEGER, D., K. MITMAN, AND F. PERRI (2016): "Macroeconomics and household heterogeneity," in *Handbook of Macroeconomics*, Elsevier, vol. 2, 843–921, https://doi.org/10.1016/bs.hesmac.2016.04.003.

KRUEGER, D., H. UHLIG, AND T. XIE (2020): "Macroeconomic dynamics and reallocation in an epidemic," *Covid Economics*, 5, 21–55, https://doi.org/10.3386/w27047.

MENDOZA, E. G., E. I. ROJAS, L. L. TESAR, AND J. ZHANG (2020): "A Macroeconomic Model of Healthcare Saturation, Inequality and the Output-Pandemia Tradeoff," Tech. rep., National Bureau of Economic Research, https://doi.org/10.3386/w28247.

Mossong, J., N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, M. Massari, S. Salmaso, G. S. Tomba, J. Wallinga, et al. (2008): "Social contacts and mixing patterns relevant to the spread of infectious diseases," *PLoS medicine*, 5, https://doi.org/10.1371/journal.pmed.0050074.

Nakajima, M. (2020): "Consumption and Saving during the Pandemic," Tech. rep., https://makotonakajima.github.io/files/paper/cscovid-paper-201030.pdf.

Piguillem, F. and L. Shi (2020): "The Optimal COVID-19 Quarantine and Testing Policies," Tech. rep., Einaudi Institute for Economics and Finance (EIEF), http://www.eief.it/eief/images/WP_20.04.pdf.

Rouse, C. and M. Gimbel (2021): "The Pandemic's Effect on Measured Wage Growth," Tech. rep., Council of Economic Advisors, https://www.whitehouse.gov/cea/blog/2021/04/19/the-pandemics-effect-on-measured-wage-growth/.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Cyclical dynamics in idiosyncratic labor market risk," *Journal of political Economy*, 112, 695–717.

U.S. Environmental Protection Agency (2020): "Mortality Risk Valuation," Tech. rep., https://www.epa.gov/environmental-economics/mortality-risk-valuation.

U.S. Federal Emergency Management Agency (2019): "Introduction to Benefit-Cost Analysis," Tech. rep., https://www.fema.gov/sites/default/files/2020-04/fema_bca_student-manual_unit-3.pdf.

Wang, H., Z. Wang, Y. Dong, R. Chang, C. Xu, X. Yu, S. Zhang, L. Tsamlag, M. Shang, J. Huang, et al. (2020): "Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China," *Cell Discovery*, 6, 1–8, https://doi.org/10.1038/s41421-020-0148-0.

A Data appendix

A.1 Details on data construction

- Inside (safe) expenditure share (Table 1): Using 2019 data on personal consumption expenditures (BEA), I compute inside consumption as the sum of expenditures on goods and housing and utilities. The inside expenditure share is then inside consumption divided by the sum of consumption expenditures on all goods and services.

- Wage ratio (Table 1): I use annual income data from the PSID core sample (2004–2014), selecting all household heads, ages 25 to 64. I compute total head labor income as the sum of the head's labor income (excluding farm and business income), head's labor part of business income, and 50 percent of household farm income, divided by two if married. Next, I construct wages by dividing head's total labor income by head's hours worked. I drop observations with missing education, with wages that are less than half of the minimum wage, with top-coded income, and with fewer than 1,000 hours per year. The ratio of the average wage for households aged 45–64 to that of households aged 25–44 is 1.35.

- Share of jobs that can be done from home (Section 3.1): Using the data provided by Dingel and Neiman (2020), I first sort occupations by their median wage. Then I place them into wage quintiles, weighted by total employment in each occupation. Finally, within each wage bin, I compute the average share of jobs that can be done from home—also provided by Dingel and Neiman (2020)—weighted by total employment.

- Per capita disposable earnings, consumption, and wealth (Table 3): I use annual income data from the PSID core sample (2014) for households whose heads are ages 25 and above. I first compute total household labor income as the sum of the head's and spouse's labor income (excluding farm and business income), head's and spouse's labor part of business income, and 50 percent of household farm income. I then use TAXSIM9 to estimate household taxes and transfers. Disposable earnings is defined as household labor income minus taxes plus transfers. Consumption is defined as expenditures on child care, clothing, education, food, health care, housing (except expenditures on mortgage, property taxes, and homeowner's and renter's insurance), transportation, vacation and entertainment, and in the case of homeowners, I add owner's equivalent rent (as in Carroll and Hur 2020b). Wealth is simply defined as household net worth

(total assets minus liabilities). I obtain per capita measures by dividing by two for married households.

A.2 Estimation of wage processes

The sample selection and estimation procedure closely follows the procedure described in Krueger et al. (2016) and Carroll and Hur (2020a). I use annual income data from the PSID core sample (1970–1997), selecting all household heads, ages 23 to 64. For waves before 1993, I use the variable Total Labor Income of Head, which is the sum of wages, tips, labor part of farm and business income, and other items. For waves after 1993, I compute total head labor income as the sum of the head's labor income (excluding farm and business income), head's labor part of business income, and 50 percent of household farm income, divided by two if married. Next, I construct wages by dividing head's total labor income by head's hours worked. I drop observations with missing education, with wages that are less than half of the minimum wage, with top-coded income, and with fewer than 1,000 hours per year. On this sample, I regress the log wage on age and education dummies, their interaction, and year dummies. I then exclude all individual wage sequences shorter than 5 years, leaving final samples of 4,524 individuals, with an average length of 9 years. On these samples, I compute the autocovariance matrix of the residuals. The stochastic process in equation (15) is estimated using GMM, targeting the covariance matrix, where the weighting matrix is the identity matrix. I thank Chris Tonetti for providing the Matlab routines that perform the estimation.

B Definition of equilibrium with transition

In this section, I provide the details regarding the equilibrium with transition. While the steady state analysis was done in general equilibrium, the transition path analysis is done in partial equilibrium. In other words, factor prices (w, r), retirement income (s), and the labor income tax (τ_ℓ) are fixed at their steady-state levels. Furthermore, the measure of newborns and their wealth distribution is also assumed to be constant throughout the transition. While these assumptions have been made for computational tractability[37], there are also other considerations. For instance, social security benefits and contribution rates do

---

[37]It is not that solving for general equilibrium dynamics would be infeasible for given a set of parameters. Calibrating the transition dynamics and solving for optimal policies in a general equilibrium setting would add considerably to the computational burden.

not typically respond to recessions at a higher frequency (and did not during the pandemic).

Even though prices are fixed on the transition path, agents still need to be given the correct path of aggregate states (Z_t), i.e. the measures of infected agents and their outside consumption and labor. This path needs to be consistent with equilibrium choices. In that sense, the transition analysis incorporates general equilibrium dynamics on the epidemiological dimension, but omits general equilibrium dynamics on the economic dimension.

A final note about the environment is that firms are not optimizing along the transition path. Instead, they rent supply-determined labor and capital to produce output and, as a consequence, suffer losses during the transition path.[38] The formal definition is provided below.

Definition. A *competitive (partial) equilibrium with transition*, given an initial distribution $\{\mu_{j1}^*\}_{j\in J}$, prices $\{w, r\}$, and fiscal policies $\{\tau_{ct}, T_t, \bar{\ell}_t^o, \tau_\ell, s\}_{t=1}^\infty$, is a sequence of value functions $\{\{v_{jt}\}_{j\in J}\}_{t=1}^\infty$, policy functions $\{\{c_{ijt}, c_{ojt}, \ell_{ijt}, \ell_{ojt}, k_{jt}'\}_{j\in J}\}_{t=1}^\infty$, producer plans $\{Y_{ft}, L_{ft}, K_{ft}\}_{t=1}^\infty$, the distribution of newborns ω, measures $\{\{\mu_{jt}\}_{j\in J}\}_{t=1}^\infty$, and aggregate states $\{Z_t\}_{t=1}^\infty$ such that, for all $t \geq 1$:

1. Given prices, fiscal policies, and aggregate states, retirees and workers solve (3) and (4), respectively.

2. Producer plans satisfy

$$Y_{ft} = K_{ft}^\alpha L_{ft}^{1-\alpha} \tag{20}$$

$$L_{ft} = \int_X \sum_{j<J_R} (\ell_{ijt}(k,\varepsilon,h) + \ell_{ojt}(k,\varepsilon,h))\, d\mu_{jt}(k,\varepsilon,h) \tag{21}$$

$$K_{ft} = \int_X \sum_{j\in J} k\, d\mu_{jt}(k,\varepsilon,h). \tag{22}$$

3. The government budget for mitigation policies clears in present value:

$$\sum_{t=1}^\infty \left\{ (1+r)^{1-t}\tau_{ct} \int_X \sum_{j\in J} (c_{ijt}(k,\varepsilon,h) + c_{ojt}(k,\varepsilon,h))\, d\mu_{jt}(k,\varepsilon,h) \right\} \tag{23}$$

$$= \sum_{t=1}^\infty \left\{ (1+r)^{1-t}T_t \int_X \mathbb{1}_{\{\ell_{ijt}(k,\varepsilon,h)+\ell_{ojt}(k,\varepsilon,h)=0\}} \sum_{j<J_R} d\mu_{jt}(k,\varepsilon,h) \right\}.$$

---

[38]Though not explicitly modeled here, the Paycheck Protection Program can be interpreted as a way to keep firms operating in this environment. An alternative approach would be to assume a small open economy environment, in which domestic capital markets do not clear. In this case, aggregate capital demanded by firms would decline, leading to a larger decline in output.

4. For any subset $(\mathcal{K}, \mathcal{E}, \mathcal{H}) \in \mathcal{B}$, the measure μ_{jt} satisfies, for $j > 1$,

$$\mu_{jt}(\mathcal{K}, \mathcal{E}, \mathcal{H}) = \int_X \psi_{j-1} \mathbb{1}_{\{k'_{j-1,t-1}(k,\varepsilon,h) \in \mathcal{K}\}} \sum_{\varepsilon' \in \mathcal{E}} \sum_{h' \in \mathcal{H}} \Gamma_{\varepsilon\varepsilon'} \Pi_{jhh't} d\mu_{j-1,t-1}(k,\varepsilon,h) \qquad (24)$$

$$+ \int_X (1-\psi_j) \mathbb{1}_{\{k'_{j,t-1}(k,\varepsilon,h) \in \mathcal{K}\}} \sum_{\varepsilon' \in \mathcal{E}} \sum_{h' \in \mathcal{H}} \Gamma_{\varepsilon\varepsilon'} \Pi_{jhh't} d\mu_{j,t-1}(k,\varepsilon,h)$$

and

$$\mu_{1t}(\mathcal{K}, \mathcal{E}, \mathcal{H}) = \int_X (1-\psi_1) \mathbb{1}_{\{k'_{1,t-1}(k,\varepsilon,h) \in \mathcal{K}\}} \sum_{\varepsilon' \in \mathcal{E}} \sum_{h' \in \mathcal{H}} \Gamma_{\varepsilon\varepsilon'} \Pi_{jhh't} d\mu_{1,t-1}(k,\varepsilon,h) \qquad (25)$$

$$+ \omega(\mathcal{K}, \mathcal{E}, \mathcal{H}).$$

5. The aggregate states $Z_t \equiv \{\mu_{It}, C^o_{It}, L^o_{It}, \epsilon_t\}$, for any $\{\epsilon_t\}^{\infty}_{t=0}$, satisfy:

$$\mu_{It} = \int_X \sum_{j \in J} \mathbb{1}_{\{h=I\}} \mu_{jt}(k,\varepsilon,h) \qquad (26)$$

$$C^o_{It} = \int_X \sum_{j \in J} \mathbb{1}_{\{h=I\}} c_{ojt}(k,\varepsilon,h) \mu_{jt}(k,\varepsilon,h) \qquad (27)$$

$$L^o_{It} = \int_X \sum_{j \in J < J_R} \mathbb{1}_{\{h=I\}} \ell_{ojt}(k,\varepsilon,h) \mu_{jt}(k,\varepsilon,h) \qquad (28)$$

C   Computational appendix

The solution algorithm broadly consists of two steps:

1. Solve for the pre-pandemic steady state.

2. Solve for pandemic transition path.

In each step, I solve the worker and retiree problems over an unevenly spaced grid of 120 wealth points, k_{coarse}. To improve solution accuracy, I place more points near zero, where the value function is more concave. I store the equilibrium wealth distribution as a histogram over an unevenly spaced wealth grid of 2000 points, k_{fine}. I set the maximum wealth level on k_{fine} much lower than the one on k_{coarse} and check that this upper bound is not overly restrictive by verifying that the equilibrium distribution has no mass on the highest grid point in the steady state or at any point along the transition.

To calibrate the pre-pandemic steady state, I guess a vector of parameters $[\beta, \varphi, \bar{u}, s, \tau_\ell]$. I then solve for the steady state, calculate the model-implied values for the targets, and update the guess using a quasi-Newton method with some dampening.

C.1 Solving for a steady state

1. Let $\mu_j^{init}(k, \varepsilon, h)$ be an initialization of the distribution over k_{fine}, E, and H.

2. Solve for the equilibrium rental rate, $r^\star$.

    (a) Guess r^0.

    (b) Given r^0, use equations (7) and (8) to get the wage,

    $$w^0(r^0) = (1 - \alpha) \left(\frac{r + \delta}{\alpha} \right)^{\frac{\alpha}{\alpha - 1}}.$$

    (c) Starting at $j = \bar{J}$, iterate on the Bellman equation in (3) until the value function converges to find the retiree value and policy functions, conditional on prices. Repeat for $j = \bar{J} - 1, .., J^R$ if $\bar{J} > J^R$.

    (d) Starting at $j = J^R - 1$, iterate on the Bellman equation in (4) until the value function converges to find the worker value and policy functions, conditional on prices. Repeat for $j = J^R - 2, .., 1$ if $J^R > 2$.

    (e) Use linear interpolation to map the policy functions from k_{coarse} onto k_{fine}.

    (f) Beginning at μ_j^{init}, update the distribution using equations (10)–(11) and the fine-grid decision rules for saving. Repeat until μ_j converges to $\mu_j^\star(r^0)$.

    (g) Use $\mu_j^\star$ and the fine-grid decision rules to compute all aggregates, including K^0 and L^0.

    (h) Find the implied interest rate, $\bar{r}^0 = \alpha \left(\frac{K^0}{L^0} \right)^{\alpha - 1} - \delta$.

    (i) Use Brent's method to solve for $r^\star$ over a fixed interval.

C.1.1 Value of statistical life

The value of $\bar{u}$ only affects the level of the value function and does not affect the policy functions in the pre-pandemic steady state. Therefore, the value of $\bar{u}$ can be found after first calibrating $[\beta, \varphi, s, \tau_\ell]$. To calculate an individual's VSL, I take the following steps:

1. I start with the pre-pandemic steady state value function:

$$\begin{aligned} v_j(k, \varepsilon, h) =& \frac{\left((c_i^*)^\gamma (c_o^*)^{1-\gamma} \right)^{1-\sigma}}{1 - \sigma} - g(\ell_i^* + \ell_o^*) + \bar{u} \\ &+ \beta \sum_{\varepsilon' \in E} \Gamma_{\varepsilon, \varepsilon'} \left[\psi_j v_{j+1}(k', \varepsilon', h) + (1 - \psi_j) v_j(k', \varepsilon', h) \right] \end{aligned}$$

where $c_i^*, c_o^*, \ell_i^*, \ell_o^*$ denote the pre-pandemic steady-state policy functions.

2. Then, imposing optimality conditions, I define

$$\hat{v}_j(k,\varepsilon,h) = \frac{\left[(c^* + \Delta_c)\,\gamma^\gamma(1-\gamma)^{1-\gamma}\right]^{1-\sigma}}{1-\sigma} - g(\ell_i^* + \ell_o^*) + \bar{u}$$
$$+ \beta(1+\Delta_s)\sum_{\varepsilon' \in E}\Gamma_{\varepsilon,\varepsilon'}\left[\psi_j v_{j+1}(k',\varepsilon',h) + (1-\psi_j)v_j(k',\varepsilon',h)\right]$$

where $c^* = c_i^* + c_o^*$ is the policy function for total consumption and Δ_c, Δ_s are small one-time deviations to consumption and survival probability.

3. Then, the VSL—defined as the marginal rate of substitution between survival and consumption—can be expressed as

$$VSL_j(k,\varepsilon,h) = \left.\frac{\frac{\partial \hat{v}}{\partial \Delta_s}}{\frac{\partial \hat{v}}{\partial \Delta_c}}\right|_{\Delta_c=0,\ \Delta_s=0}$$
$$= \frac{\beta \sum_{\varepsilon' \in E}\Gamma_{\varepsilon,\varepsilon'}\left[\psi_j v_{j+1}(k',\varepsilon',h) + (1-\psi_j)v_j(k',\varepsilon',h)\right]}{(c^*)^{-\sigma}\left(\gamma^\gamma(1-\gamma)^{1-\gamma}\right)^{1-\sigma}}$$

4. Set $\bar{u}$ such that

$$\int_X \sum_{j \in J} VSL_j(k,\varepsilon,h)d\mu_j(k,\varepsilon,h) = 4096\int_X \sum_{j \in J}\left(c_{ij}(k,\varepsilon,h) + c_{oj}(k,\varepsilon,h)\right)d\mu_j(k,\varepsilon,h).$$

C.2 Solving for a transition path

Recall that all individuals are assumed to be in the recovered state by $t = \hat{t}$ (e.g. vaccine and cure). Combined with the assumption that prices are fixed at their initial steady state values throughout the transition path, this implies that, even though the measure μ_{jt} takes a very long time to return to its steady state values, value and policy functions only need to be solved for $t \le \hat{t}$ since all the relevant variables for solving the individual's optimization problem are constant for $t > \hat{t}$ (at steady-state levels).

To introduce COVID-19 in the economy, set $\mu_{j1}(k,\varepsilon,I) = 0.005\mu_{j0}(k,\varepsilon,S)$, $\mu_{j1}(k,\varepsilon,S) = 0.995\mu_{j0}(k,\varepsilon,S)$, and $\mu_{j1}(k,\varepsilon,R) = 0$ for $j \in J$, $k \in K$, and $\varepsilon \in E$. Recall that all agents are susceptible in the pre-pandemic steady state ($t = 0$).

1. Guess the sequence $\{Z_t, \tau_{ct}\}_{t=1}^{\hat{t}}$.

2. Set $v_{j,\hat{t}+1}$ equal to the steady-state value function for $j \in J$. Then, starting in period $\hat{t}$, solve the Bellman equations in (3) and (4) backward. This produces a sequence of policy functions for periods $t = 1,..,\hat{t}$.

3. Starting at μ_{j1}, simulate forward using the policy functions to find the sequence of measures from $t = 1, ..., t_{end}$ for a very large t_{end} (recall that the stated assumptions guarantee that the policy functions are constant for $t \geq \hat{t}$). Along the way, solve for aggregate variables in each period, including $\bar{Z}_t \equiv \{\mu_{It}, C^o_{It}, L^o_{It}, \epsilon_t\}$, where ϵ_t is set so that the model generates the same times series of deaths as in the data for $t < 26$ (this step is only required for the calibration described in Section 3.2).

4. Check that the difference between the guess for $\{Z_t\}$ and the implied value $\{\bar{Z}_t\}$ (measured under the sup norm) is less than a small tolerance. Additionally, if $T > 0$, check that the government budget for mitigation policies in (23) clears in present value. If so, a transition path has been found.

5. If not, update the guess using a dampening method and repeat.

Note that, while the solution for $\{Z_t\}$ is not guaranteed to have a unique solution, I have checked that starting the routine with alternative guesses converge to the same solution.

D   Robustness

In this section, I explore the robustness of the results to alternative parameter and modeling choices. In all of the robustness exercises, I recalibrate the model to match the same moments as in the baseline calibration, with the exception of the cases in which the stay-at-home subsidy is financed with alternative tax instruments (these counterfactual exercises do not affect the calibration). I find that although the level of average welfare, deaths, and output can vary substantially, the main policy implications are extremely robust. Under all alternative configurations, the output-maximizing policy simultaneously improves output and reduces deaths, and is Pareto improving. Moreover, neither the output-maximizing nor the best Pareto-improving policies involve a lockdown. The policy parameters and selected outcomes of the best Pareto-improving and output-maximizing policies under alternative parameter values and modeling choices are summarized in Tables 6–7.

Higher value of statistical life. In the baseline calibration, I use a VSL of \$6.9 million in 2019, which corresponds to 4,096 biweekly consumption per capita. As a robustness check, I use an higher VSL recommended by U.S. Environmental Protection Agency (2020) of \$7.4 million in 2006, or 6,226 times biweekly consumption per capita. By comparing the no public mitigation results in Tables 5 and 6, I find that assigning a higher value of

statistical life leads to slightly lower deaths and output and a larger welfare loss as result of the pandemic. The main policy implications remain the same: Both the output-maximizing and best Pareto-improving policies simultaneously save lives and increase output relative to no public mitigation, and feature no lockdowns.

Smaller efficiency loss during infection.   Next, I consider a 30 percent reduction in labor efficiency during infection, compared with the baseline calibration of a 50 percent reduction. Interestingly, comparing the no public mitigation scenarios under a 30 percent versus a 50 percent efficiency reduction during infection reveals that welfare and output is actually lower under the 30 percent reduction. This is because the smaller efficiency loss for infected individuals induce these infected individuals to work more, leading to a more severe pandemic and a larger decline in economic activity among susceptible agents. The main policy implications are nevertheless unchanged.

Lower disutility of not working.   The baseline calibration ($\tilde{u} = 0.82$) induces a 15 percent decline in employment that is consistent with data (2020, *Bureau of Labor Statistics*). As a robustness check, here I consider a lower disutility of $\tilde{u} = 0.61$, which generates a 19 percent decline in employment, consistent with the real-time survey data reported by Bick et al. (2020). As can be expected, the no-mitigation case is identical to the baseline. The main policy implications are similar to the baseline, except that the subsidies are smaller in size since in this specification individuals are less averse to accepting the subsidies.

Higher basic reproduction number.   I also investigate the implications of assuming a higher basic reproduction number of 2.5, compared to the baseline value of 2.2. As expected, a higher basic reproduction number leads to a larger number of deaths. Nevertheless, the policy implications are similar to the baseline.

Higher inside consumption.   In the baseline calibration, I set the inside consumption preference parameter $\gamma = 0.51$ to match the expenditure share on goods and housing and utility services in 2019. As I discuss in the main text, an exact mapping between inside (or safe) consumption in the model to the data is not possible due to data limitations. Thus, as a robustness check, I consider a higher value of $\gamma = 0.61$. Table 6 shows that the results are nearly identical to those in the baseline.

	subsidy amount ($/week)	duration (months)	lock-down	average welfare change (percent)	2-year output (index)	deaths (per 10k)
higher value of statistical life ($VSL = 6226\bar{c}$)						
best Pareto-improving*	900	9	no	−5.1	97.2	13.5
output-maximizing*	850	11	no	−5.1	97.4	13.6
no public mitigation	0	0	no	−7.1	96.2	20.1
smaller efficiency loss during infection ($\eta_{jI} = 0.7\eta_{jS}$)						
best Pareto-improving*	700	10	no	−4.1	97.4	16.0
output-maximizing*	650	11	no	−4.2	97.4	16.4
no public mitigation	0	0	no	−6.0	96.5	24.8
lower disutility of not working ($\tilde{u} = 0.61$)						
best Pareto-improving*	600	7	no	−3.6	97.4	13.6
output-maximizing*	350	13	no	−4.4	98.3	17.4
no public mitigation	0	0	no	−5.4	97.3	22.0
higher basic reproduction number ($R_0 = 2.5$)						
best Pareto-improving*	700	12	no	−4.7	97.1	18.2
output-maximizing*	650	11	no	−4.9	97.3	19.1
no public mitigation	0	0	no	−6.7	96.1	27.8
higher inside consumption ($\gamma = 0.61$)						
best Pareto-improving*	700	9	no	−3.6	97.8	14.1
output-maximizing*	500	9	no	−4.0	98.0	16.0
no public mitigation	0	0	no	−5.3	97.3	22.1

Notes: Average welfare change reports the population-weighted average of individual consumption equivalents. Output refers to output from $t = 1$ (March 27, 2020) to $t = 52$ (March 24, 2022), compared with the analogous 52-period output in the steady state, indexed at 100. * denotes Pareto improvements relative to the no public mitigation case.

No hospital capacity constraint. I also investigate the implications of assuming that there is no hospital capacity constraint. As expected, this leads to fewer deaths and a smaller welfare loss under the no public mitigation scenario, compared with the baseline calibration. Nevertheless, the policy prescriptions are similar: the output-maximizing policy parameters are unchanged, while the best Pareto-improving policy features a smaller but slightly longer subsidy than in the baseline calibration.

Subsidy funded by tax on outside consumption. In the baseline calibration, I assumed that the stay-at-home subsidy is funded by a tax on total consumption. Here, I consider a Pigouvian tax that specifically targets outside consumption. The main policy implications are similar to that in the baseline. The most significant difference is that the Pigouvian tax is much more effective at reducing deaths. A nearly identical subsidy brings down deaths to 10.7 (per 10 thousand), compared to 14.0 in the baseline. This is because the Pigouvian tax more directly reduces outside consumption, leading to less infections and deaths.

Subsidy funded by tax on labor income. Here, I use a labor income tax to fund the stay-at-home subsidy. Compared to the baslinee, the policy prescriptions are similar except with smaller subsidies, resulting in smaller increases in output and smaller reductions in deaths.

Deceased individual replaced by newborn. The baseline model assumes that accidental bequests from deceased individuals are given to a fraction of newborn individuals each period. Here, I consider an alternative assumption: a deceased individual is replaced by a newborn with the same productivity and wealth level. I find that the main policy prescriptions are similar to the baseline model.

No pandemic. Given the relatively large, broad, and robust support of the mitigation policies, one might wonder whether such policies would be beneficial in normal times. I find that this is not the case. In fact, in the absence of a pandemic, there is no policy configuration that delivers higher output or an average welfare gain, let alone one with unanimous support.

Table 7: Sensitivity analysis (2)

	subsidy amount ($/week)	duration (months)	lock-down	average welfare change (percent)	2-year output (index)	deaths (per 10k)
no hospital capacity constraint ($\kappa = 1$)						
best Pareto-improving*	550	11	no	–4.1	97.9	15.8
output-maximizing*	500	10	no	–4.2	98.0	16.5
no public mitigation	0	0	no	–5.3	97.4	21.3
subsidy funded by tax on outside consumption						
best Pareto-improving*	700	11	no	–2.9	98.0	10.7
output-maximizing*	500	10	no	–3.7	98.2	14.4
no public mitigation	0	0	no	–5.4	97.3	22.0
subsidy funded by tax on labor income						
best Pareto-improving*	500	10	no	–4.1	97.5	16.1
output-maximizing*	400	12	no	–4.7	97.8	18.9
no public mitigation	0	0	no	–5.4	97.3	22.0
deceased individual replaced by newborn						
best Pareto-improving*	600	11	no	–3.7	97.9	14.5
output-maximizing*	400	10	no	–4.5	98.0	18.6
no public mitigation	0	0	no	–5.2	97.4	21.7
no pandemic						
best Pareto-improving	0	0	no	0.0	100.0	0.0
output-maximizing	0	0	no	0.0	100.0	0.0
no public mitigation	0	0	no	0.0	100.0	0.0

Notes: Average welfare change reports the population-weighted average of individual consumption equivalents. Output refers to output from $t = 1$ (March 27, 2020) to $t = 52$ (March 24, 2022), compared with the analogous 52-period output in the steady state, indexed at 100. * denotes Pareto improvements relative to the no public mitigation case.