Regular article

The macroeconomics of pandemics around the world: Lives versus livelihoods revisited[☆]

Ingvild Almås [a,b], Tessa Bold [a,*], Tillmann von Carnap [a], Selene Ghisolfi [c], Justin Sandefur [d]

[a] *Institute for International Economic Studies, Stockholm University, Sweden*
[b] *Norwegian School of Economics, Bergen, Norway*
[c] *Università Cattolica del Sacro Cuore, Milan and LEAP, Bocconi University, Milan, Italy*
[d] *Center for Global Development, Washington D.C., United States of America*

ARTICLE INFO

ABSTRACT

The COVID-19 pandemic led governments around the world to impose unprecedented restrictions on economic activity. Were these restrictions equally justified in poorer countries with fewer demographic risk factors and less ability to weather economic shocks? We develop and estimate a fully specified model of the macroeconomy with epidemiological dynamics, incorporating subsistence constraints in consumption and allowing preferences over "lives versus livelihoods" to vary with income. Poorer countries' demography pushes them unambiguously toward laxer policies. But because both infected and susceptible agents near the subsistence constraint will remain economically active in the face of infection risk and even to some extent under government containment policies, optimal policy in poorer countries pushes in the opposite direction. Moreover, for reasonable income-elasticities of the value of a statistical life, the model can fully rationalize equally strict or stricter policies in poorer countries.

1. Introduction

The COVID-19 pandemic led governments around the world to impose unprecedented restrictions on economic activity. These measures were surprisingly uniform across countries at all income levels: throughout 2020, low-income countries enacted policies roughly as stringent as those in high-income countries (Fig. 1). In the United States, a survey in late March 2020 found zero leading economists disagreed that the policy response to the pandemic should involve "a very large contraction in economic activity until the spread of infections has dropped significantly" (IGM Forum, 2020). At the same time, many economists expressed reservations about applying similar policy prescriptions to developing countries (Ray et al., 2020; Ray and Subramanian, 2020; Barnett-Howell et al., 2021; Ravallion, 2020; Miguel and Mobarak, 2022).

In this paper, we ask whether poor countries were justified in imposing equally harsh restrictions as rich countries – despite fewer demographic risk factors and lower ability to weather economic shocks. We show that these similar policy responses between rich and poor countries can be rationalized by extending a standard macroeconomic model, with feedback between the economy and the pandemic, to allow for the trade-off between lives and livelihoods to vary by income.[1] Empirical patterns are consistent with the model predictions under two, very reasonable, conditions: (a) poor countries are close to their subsistence constraint below which consumption cannot fall, and (b) the valuation of lives versus income is increasing in income which means that at the margin, poor countries value lives relative to income more than rich countries.

Our analysis starts from the model of Eichenbaum et al. (2021), henceforth ERT, in which agents expose themselves to infection risk
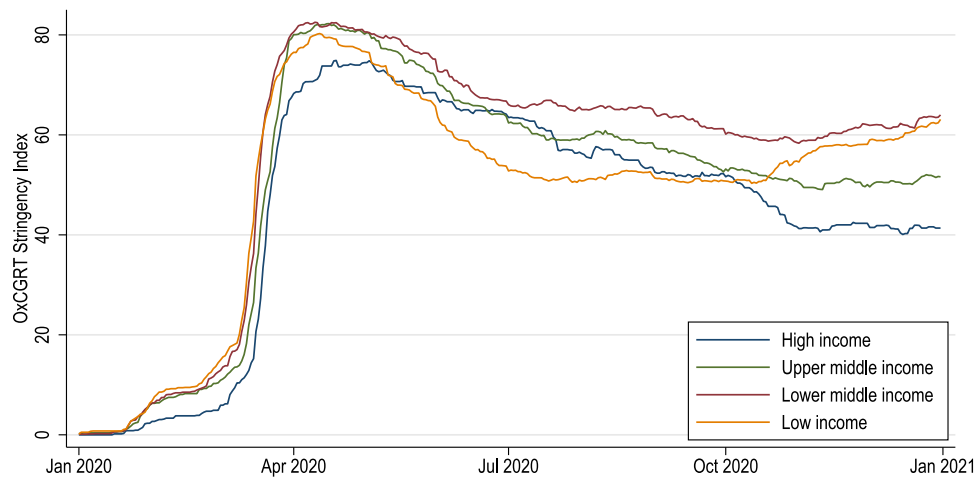
Fig. 1. Stringency of pandemic containment policies during 2020 by income group. The figure shows the daily average in the level of restrictions as measured by the Oxford Coronavirus Government Response Stringency Index (Hale et al., 2021) across countries in income groups as defined by the World Bank. The stringency index varies between 0 and 100 and calculates an average of all containment measures a country has taken, including school closures, workplace closures, cancellation of public events, restrictions on gathering size, closure of public transport, and stay-at-home requirements.

when working and consuming, and, realizing this danger, reduce economic activity as infection risk rises. However, they do not weigh the impact of their own labor and consumption decisions on the pandemic's spread, creating an externality which the social planner seeks to internalize. The model allows for a policy lever in the form of a value-added tax on consumption, discouraging economic activity and hence slowing disease transmission. The policymaker chooses the tax to maximize the present value of aggregate utility, taking into account utility losses from both reduced consumption and lost lives.

We generalize ERT's model by incorporating two features that may drive differences in optimal policy between countries at different income levels and with different fatality risk from COVID-19: a subsistence constraint, and varying values of a statistical life (VSL). First, to capture the intuition that poorer people may be less able to reduce their economic activity during the pandemic than richer counterparts, we introduce a subsistence constraint below which consumption cannot fall. Second, we also allow the trade-off between consumption and catching a potentially deadly infection to vary by income. That is, the VSL will always be lower in poorer countries, but the VSL-income ratio may not. Specifically, we allow the VSL to be a non-linear function of income and show that optimal policy depends crucially on the income-elasticity of the VSL, as the social planner trades off risks of deaths against the utility of the living. We estimate both these parameters, i.e. the subsistence constraint and the income-elasticity of the VSL, using indirect inference.

The model highlights two important points. First, recognizing that foregoing short-term consumption may be a question of survival for poor households has somewhat counter-intuitive implications for policy. Because both infected and susceptible agents, who are pushed close to their subsistence constraint will remain economically active in the face of infection risk and even government containment policies, optimal policy in poorer countries becomes more rather than less strict. Second, comparative statics illustrate that any statement about the desirability of laxer lockdowns in poorer countries is highly sensitive to the VSL and how it changes with income.

Next, we examine how pandemic outcomes and policies varied by income and infection fatality rates (IFR) for 125 countries through the lens of this model. We find, for instance, that, relative to rich countries, poor countries had slightly stricter containment policies during the first year of the pandemic (i.e. before vaccines arrived), but nevertheless experienced smaller drops in economic activity. They also saw higher infections yet fewer deaths. On the other hand, countries with low IFR display laxer containment policies, lower economic contractions, and less excess mortality despite higher infection rates.

We then use indirect inference to estimate the unobserved subsistence constraint and income elasticity of the VSL that produce a close match between data and model in the difference in pandemic outcomes across rich and poor as well as high- and low IFR countries. The model successfully matches the data with an income-elasticity of the VSL which is strictly less than unity – implying the value countries place on lost lives increases less than proportionally with income – and a subsistence constraint that is strictly greater than zero. The value for the income-elasticity of the VSL is in line with estimates in Viscusi and Masterman (2017) who suggest an income-VSL elasticity between 0.8 and 1.2, and lines up with recent cross-country estimates of the income-elasticity of health expenses (Farag et al., 2012). A positive subsistence level is consistent with a lower elasticity of substitution between labor and consumption at low income levels (see King et al. (1988)) that is commonly found in historical data (Ohanian et al., 2008; Bick et al., 2018; Boppart and Krusell, 2020).

A key contribution of our analysis is to show that a relatively parsimonious SIRD (Susceptible–Infected–Recovered–Deceased) macro model can explain broad qualitative patterns of the aggregate variables and the policy responses across countries. In terms of the model's implications, our estimation implies that relatively strict containment policies by developing countries in the early stages of the pandemic can be rationalized if policymakers in the developing world placed a relatively higher value on life (though well within the existing range of empirical estimates) and households faced subsistence constraints which limited their spontaneous social distancing and their compliance with containment policies.

Our study contributes to an emerging literature assessing the welfare implications of pandemic containment policies in low- and high-income countries using macroeconomic models (Alon et al., 2020; Hausmann and Schetter, 2022). Alon et al. (2020) assume that workers in the informal sector cannot be shielded from the disease by a lockdown and argue that containment policy will hence be less effective in countries with larger informal sectors. Our approach highlights a similar effect, though grounded in the utility maximization of agents which in turn affects the planner's optimal policy: when faced with a given risk of contracting the disease through economic activity, poorer agents will rationally reduce their exposure less, requiring relatively stricter policies to achieve the same reduction in deaths. Alon et al. (2020) further emphasize that demographic differences, as captured by the country-varying IFR in our framework, account for most of the differences in mortality rates between their modeled rich and poor countries. Our approach illustrates that lower mortality risk may

not only mechanically affect the overall death rate, but also influence individual-level optimization and adaptation behavior. In contrast to Hausmann and Schetter (2022)'s two-period model of households facing subsistence constraints, we jointly model the full paths of infections and optimal policy and allow for households' decisions to endogenously affect the pandemic's development.

Our paper further complements studies explaining different pandemic and economic outcomes by levels of incomes both within and across countries (Eichenbaum et al., 2022; Kim et al., 2023). Similar to our study, Eichenbaum et al. (2022) introduce different VSLs and a subsistence constraint to understand heterogeneous behavioral reactions and health outcomes among different income groups within the United States. Our paper differs in its focus on developing countries, and validates key model choices for optimal policy determination: the subsistence constraint and the scaling parameter for the VSL. We share the cross-country perspective with Kim et al. (2023), but differ in our focus on individual-level behaviors affecting aggregate optimal policy, as opposed to explaining cross-country outcomes.

The macro-pandemic model we consider adds a global perspective to recent work incorporating economic decision-making into the SIRD framework. According to these integrated models, agents facing contagion risk will voluntarily reduce their economic activity, thus partly containing the virus (Toxvaerd, 2020; Garibaldi et al., 2020; Chudik et al., 2022). However, analyses suggest that, from a social welfare perspective, further government action is justified by agents' failure to internalize their own contribution to the spread of infections (Eichenbaum et al., 2021; Farboodi et al., 2020; Krueger et al., 2022; Glover et al., 2021; Alvarez et al., 2021).

In the following section, we lay out our extensions to the basic model, combining epidemiological and economic frameworks. We calibrate the model for two pairs of representative countries – rich and poor, and high and low IFR – and present comparative statics across income, IFR, and preferences for lives vs. income in Section 3. We present empirical estimates of the evolution of excess mortality, infections, lockdowns and economic contraction in Section 4. We structurally estimate and validate the model in Section 5. Section 6 discusses the plausibility of the structural estimates and Section 7 concludes with implications for policy.

2. Model

2.1. The pandemic

We begin with a standard SIRD model to describe the evolution of the pandemic over time. The model compartmentalizes the population into four groups, namely susceptibles, S_t, who have not yet had the disease, infected, I_t, who are currently sick and can infect others, those who have already recovered (and are assumed to have acquired lasting immunity) R_t, and deceased, D_t. The stocks of these groups in the population are described by the following system of difference equations. First, the number of susceptibles in period $t + 1$ equals the stock in the previous period minus the number of newly infected in period t, T_t:

$$S_{t+1} = S_t - T_t$$

Second, the number of infected in period $t + 1$ equals the stock of infected in period t plus those who are newly infected in period t minus those who either recovered in the previous period (share π_r of the infected) or died (share π_d of the infected). The probability π_d is the infection fatality rate (IFR), i.e. the probability of dying from COVID-19 conditional on being infected:

$$I_{t+1} = I_t + T_t - (\pi_r + \pi_d)I_t$$

Third, the stock of the recovered at time $t + 1$ equals the stock of recovered in period t plus those who recovered in period t:

$$R_{t+1} = R_t + \pi_r I_t$$

Finally, the stock of the deceased increases by those who were infected and died in the previous period:

$$D_{t+1} = D_t + \pi_d I_t$$

The population therefore evolves according to

$$Pop_{t+1} = Pop_t - \pi_d I_t.$$

The pandemic ends once a sufficiently large share of the living population has acquired immunity. In Section 3, we detail how we calibrate the transition probabilities between the different states.

2.2. The economy during the pandemic

We generalize the model of Eichenbaum et al. (2021) by incorporating three features that may drive differences in optimal policy between countries at different income levels. First, we allow poorer agents to respond less to the pandemic threat and any containment measures to combat it. We achieve this by introducing a subsistence constraint which is standard in demand system estimations going back to Stone (1954). Such a constraint has been shown to be consistent with historical data on consumption and leisure choices (Bick et al., 2018) as well as with evidence from the United States during the early stages of the pandemic showing that richer households accounted for the bulk of spending reductions (Chetty et al., 2021). Second, when calibrating the model in later sections, we allow two key characteristics of a country that matter for its pandemic response – IFR and income – to vary across contexts. Third, we make the (otherwise implicit) value of statistical life an explicit function of income.

The economy consists of representative agents for each of the susceptible, the infected and the recovered groups. Each type of agent, $j = s, i, r$, maximizes the discounted sum of instantaneous utility u, which is a function of their consumption c_t^j and hours worked n_t^j:

$$U^j = \sum_{t=0}^{\infty} \beta^t u(c_t^j, n_t^j) \tag{1}$$

The instantaneous utility function u depends on two additional parameters: (i) we allow for a subsistence constraint for consumption by including a constant $\bar{c}$ that may or may not be zero. In general, this constant is thought of as non-negative and representing a consumption level below which income cannot fall. With a positive subsistence constraint the elasticity of consumption with respect to a tax is increasing in income. (ii) We target specific VSLs by incorporating a baseline utility value of being alive, $\bar{u}$, which is increasing in the targeted VSL. The target VSL is an increasing function of income, and we specify below how it is obtained. With log-preferences for consumption and a convex disutility of labor, this gives

$$u(c_t^j, n_t^j) = \ln(c_t^j - \bar{c}) - \frac{\theta}{2}(n_t^j)^2 + \bar{u}. \tag{2}$$

The budget constraint for agent $j = s, i, r$ is given by

$$(1 + \mu_{ct})c_t = w_t \phi_t^j n_t + \Gamma_t. \tag{3}$$

where w_t is the wage, μ_{ct} is a value added tax that discourages consumption – the social planner's policy lever that we explain below – and Γ_t is a lump-sum transfer from the government. ϕ_t^j is a productivity parameter that equals 1 for the susceptible and recovered and less than 1 for the infected.

Output in the economy is generated by a representative firm, which produces a consumption good using labor with the technology $C_t = AN_t$. Hours worked are chosen to maximize firm's profits $\Pi_t = AN_t - w_t N_t$ with the first-order condition $w_t = A$. The government's budget constraint is

$$\mu_{ct}(c_t^s S_t + c_t^i I_t + c_t^r R_t) = \Gamma_t(S_t + I_t + R_t). \tag{4}$$

The model is closed by the equilibrium conditions that agents' labor supply equals the firm's labor demand, $n_t^s S_t + \phi^i n_t^i I_t + n_t^r R_t = N_t$ and that their consumption equals the output produced, $c_t^s S_t + c_t^i I_t + c_t^r R_t = C_t$.

We now explain how the economy and the pandemic interact. In the standard SIRD model, the probability of infection is given by a single parameter independent of agents' behavior. Following ERT, we instead allow economic activity to raise the risk of infection because susceptible and infected interact with each other in order to consume and work. The probability that a susceptible agent becomes infected through consumption activities is assumed to be proportional to their own consumption, c_t^s, and the total consumption of the infected, $(I_t c_t^I)$. It is denoted by $\pi_{s1} c_t^s (I_t c_t^I)$. Similarly, the probability that a susceptible becomes infected at work is proportional to their hours worked, n_t^S, and the total hours worked by the infected $I_t n_t^I$. It is denoted by $\pi_{s2} n_t^S (I_t n_t^I)$. Finally, just as in the standard SIRD, we assume that some infections happen irrespective of the intensity of economic interactions, and the probability that a susceptible contracts the disease in this way is simply proportional to the number of infected, i.e. $\pi_{s3} I_t$.

With these assumptions, the probability that a susceptible becomes infected in period t via one of the three channels is

$$\tau_t = \pi_{s1} c_t^S (I_t C_t^I) + \pi_{s2} n_t^S (I_t N_t^I) + \pi_{s3} I_t. \tag{5}$$

The proportionality constants π_s ensure that τ_t as well as each of its the three components are well-defined probabilities. To give an intuitive interpretation of these parameters, π_{s1} (π_{s2}) is the probability that a susceptible gets infected when they meet another infected and they both consume (work to earn) one dollar. It would be larger, for instance, the more contagious is the disease, the more time is spent shopping (working) per each dollar consumed (earned), or the higher the interaction between susceptibles and infected during consumption (work) activities.[2]

The endogeneity of infection risk creates a feedback loop between the economy and the pandemic: susceptible agents want to avoid infection both because it lowers their productivity temporarily and because it carries a risk of dying and thus foregoing utility from consuming and being alive. As a consequence, they reduce consumption and hours worked, which in turn slows the spread of the pandemic. Specifically, the discounted lifetime utility of a susceptible person is

$$U_t^s = u(c_t^s, n_t^s) + \beta[(1 - \tau_t)U_{t+1}^s + \tau_t U_{t+1}^i],$$

where U_{t+1}^i is the discounted lifetime utility when infected. This, in turn, is given by

$$U_t^i = u(c_t^i, n_t^i) + \beta[(1 - \pi_r - \pi_d)U_{t+1}^i + \pi_r U_{t+1}^r],$$

where $(1 - \pi_r - \pi_d)$ is the probability of still being infected in period $t+1$, π_r is the probability of recovering in period $t+1$ and U_{t+1}^r is the lifetime utility after recovery. Thus, when the baseline utility $\bar{u} = 0$, the cost of dying, which happens with probability π_d when infected, equals the discounted lifetime utility of consumption and leisure the individual must forego. Finally, the discounted lifetime utility of recovered agents is simply

$$U_t^r = u(c_t^r, n_t^r) + \beta U_{t+1}^r.$$

We can find optimal consumption and hours worked by maximizing each agent's $j = s, i, r$ discounted lifetime utility subject to their budget constraint by way of a Lagrangian. For the susceptible the maximization involves an additional choice, namely the optimal value for the probability of becoming infected τ_t, and an additional constraint, namely the dependence of this probability on the agent's consumption and labor supply (Eq. (5)). Letting λ_{bt}^j denote the Lagrange multiplier on each agent's budget constraint and λ_τ the Lagrange multiplier on the infection probability constraint, the constrained maximization results in the

following first-order conditions for consumption, $c_t^{s,i,r}$ (Eqs. S, I, R (i)), hours, $n_t^{s,i,r}$ (Eqs. S, I, R (ii)), and (for the susceptible) the endogenous probability of infection, τ_t (Eq. $S(iii)$):

$$
\begin{array}{llll}
\text{S:} & (i) & u_1\left(c_t^s, n_t^s\right) - \lambda_{bt}^s \left(1 + \mu_{ct}\right) + \lambda_{\tau t}\pi_{s1}\left(I_t c_t^I\right) & = 0 \\
& (ii) & u_2\left(c_t^s, n_t^s\right) + \lambda_{bt}^s w_t + \lambda_{\tau t}\pi_{s2}\left(I_t n_t^I\right) & = 0 \\
& (iii) & \beta(U_{t+1}^i - U_{t+1}^s) - \lambda_{\tau t} & = 0 \\
\text{I:} & (i) & u_1\left(c_t^i, n_t^i\right) - \lambda_{bt}^i \left(1 + \mu_{ct}\right) & = 0 \\
& (ii) & u_2\left(c_t^i, n_t^i\right) + \lambda_{bt}^i \phi^i w_t & = 0 \\
\text{R:} & (i) & u_1\left(c_t^r, n_t^r\right) - \lambda_{bt}^r \left(1 + \mu_{ct}\right) & = 0 \\
& (ii) & u_2\left(c_t^r, n_t^r\right) + \lambda_{bt}^r w_t & = 0
\end{array}
$$

The additional terms in the first-order conditions for the susceptible ($\lambda_{\tau t}\pi_{s1}\left(I_t c_t^I\right)$ for consumption and $\lambda_{\tau t}\pi_{s2}\left(I_t n_t^I\right)$ for hours worked) imply that susceptible agents respond to the pandemic by reducing hours and consumption even in the absence of any containment measures. In general, this endogenous response will reduce economic activity less than what is socially optimal, because infected agents fail to internalize the risk of exposing others. There is thus a role for government to use taxes to reduce economic activity and contain the pandemic further. The optimal sequence of containment rates during the course of the pandemic $\{\mu_{ct}\}_{t=0}^T$ is chosen by the social planner to maximize the sum of lifetime utilities of the three types of agents in the economy, subject to the government's budget constraint. That is, the social planner solves

$$\max_{\{\mu_{ct}\}_{t=0}^T} (S_0 U_0^s + I_0 U_0^i), \tag{6}$$

subject to

$$\mu_{ct}(c_t^s S_t + c_t^i I_t + c_t^r R_t) = \Gamma_t(S_t + I_t + R_t) \quad \forall t,$$

where agents' choices, in turn, are optimal given the containment rate.[3]

3. Model calibration

We now analyze the model's predictions via simulations. In a first step, we illustrate its main mechanisms for a benchmark high-income country, the United States. In a second step, we examine how optimal containment policy and agents' response to it change under scenarios with differing levels of disease risk and poverty. In both cases, the model is set to a weekly frequency.

3.1. Benchmark calibration

For our benchmark calibration, we set the value of the newly introduced parameters, $\bar{c}$ and $\bar{u}$, to zero. We describe the calibration of the remaining parameters below and in column 2 of Table 1.

To calibrate the economic part of the model, we follow ERT in setting the weekly discount factor to $\delta = 0.96^{\frac{1}{52}}$. We calibrate labor productivity A, and the disutility of labor θ using information on United States' pre-pandemic income and hours worked: we convert World Bank annual GDP figures at 2010 constant dollars into weekly values, $Y_{weekly,US} = \frac{55,753}{52}$ and set steady-state weekly hours worked for the United States, $H_{weekly,US}$, to 30 (Charmes, 2015). Using the representative firm's production function, we calculate the labor productivity parameter for the United States as $A_{US} = \frac{Y_{weekly,US}}{H_{weekly,US}}$. From this, we can calculate steady-state weekly consumption as $c_{weekly,US}^* = \frac{A}{(H_{weekly,US})}$. We find the disutility of labor parameter for the United States by inverting the household's first-order condition for hours worked in the absence of a pandemic: $\theta_{US} = \frac{1}{(H_{weekly,US})^2}$.

---

[2] Having defined the probability of infection in period t, τ_t, the sum of new infections from the three channels in period t equals:

$$T_t = \tau_t S_t = \pi_{s1}(c_t^S)(I_t c_t^I) + \pi_{s2}(n_t^S)(I_t n_t^I) + \pi_{s3}(I_t S_t).$$

[3] Note that there are only terms for the susceptible and the infected in the planner's objective function because there are no recovered agents at the beginning of the pandemic.

Table 1
Mapping model parameters to data.

	Definition	Benchmark	Representative countries	All countries	Source
Estimated parameters					
$\bar{c}$	Subsistence constraint (if < 0: home production)	0	$\bar{c} \in [-200, 600]$, precise value to be estimated	$\bar{c} \in [-200, 600]$, precise value to be estimated	
η	Elasticity of VSL to income	NA	$\eta \in [0.7, 1.2]$, precise value to be estimated	$\eta \in [0.7, 1.2]$, precise value to be estimated	Viscusi and Masterman (2017)
Calibrated parameters					
β	Annual discount factor	0.96	0.96	0.96	ERT
A	Productivity of labor	35.74	Calculated using steady-state weekly hours in representative countries	From country-specific weekly hours	ERT, World Development Indicators; Charmes (2015)
θ	Disutility of labor	0.0011	Calculated using steady-state weekly hours in representative countries	From country-specific weekly hours and GDP/capita	ERT; Charmes (2015)
ϕ	Productivity while infected	0.8	0.8	0.8	ERT
π_d	Prob. of dying within 7 days of infection	$0.0064 \times \frac{7}{18}$	IFR of representative country	Country-specific predicted IFR	ERT; Ghisolfi et al. (2020)
π_r	Prob. of recovering within 7 days of infection	$(1 - 0.0064) \times \frac{7}{18}$	,,	,,	
π_{sk}	Scaling parameters in a susceptible's probability of getting infected if exposed during - consumption ($k = 1$) - work ($k = 2$) - other contacts ($k = 3$)	$\pi_{s1} = 7.84 \times 10^{-8}$ $\pi_{s2} = 1.24 \times 10^{-4}$ $\pi_{s3} = 0.3901$	Calculated using steady-state consumption and hours in representative countries	Calculated using steady-state consumption and hours in each country	ERT

To calibrate the pandemic part of the model, we follow ERT in setting the productivity loss from getting infected at 20% (i.e $\phi = 0.8$), and also in assuming that the average time to either recovery or death is 18 days.

For the probability of dying from COVID-19 conditional on being infected, we depart from ERT and use country-specific predictions from Ghisolfi et al. (2020), where a country's IFR is predicted as a function of its demography (younger populations have a lower IFR), the prevalence of relevant comorbidities, and its health system capacity (health systems with a proven track record of treating viral respiratory diseases are assumed to exhibit a lower IFR). For the United States, the authors predict an IFR of 0.64 percent. This results in a weekly probability of dying from COVID-19 when infected of $\pi_d = \frac{7}{18} \times 0.0064$ and a weekly probability of recovering of $\pi_r = \frac{7}{18}(1 - 0.0064)$ for the benchmark case.

Finally, we follow ERT in calibrating the parameters π_{s1}, π_{s2}, π_{s3} in the expression for the transmission probability of COVID-19 (Eq. (5)). Specifically, we set the initial share of the population that is infected to 0.001 and simulate the evolution of the pandemic under the assumption that the susceptible and infected work steady-state hours and have steady-state consumption in each period for an initial guess of the π_{sk}. We then adjust the π_{sk} until 1/6 of infections each period are due to work interactions, 1/6 are due to consumption interactions and 2/3 are due to random interactions, and 60% of the population have been infected by week 250 of the pandemic.

Given the calibrated π_{sk}, the share of initial infections and steady-state hours and consumption, we can calculate the average number of secondary infections produced by a typical case of an infection in a population where everyone is susceptible, which is basic reproduction number R_0 of the pandemic (see Eichenbaum et al. (2021), p.5163). For the US benchmark case, this amounts to 1.5, which is at the lower end of estimates reported for the early stage of the pandemic in Liu et al. (2020).

Fig. 2 shows the results of our benchmark calibration. The dotted black line displays the course of the epidemic without any behavioral adjustment or government intervention — the basic SIRD model as outlined in Section 2.1.[4] The epidemic cannot end before a sufficient share of the population has acquired immunity. Given a basic reproduction number of $R_0 = 1.5$, the herd immunity level is 33%.[5] However, the rapid increase in infections in the pure SIRD-model leads to an overshoot of infections (and deaths) above this herd-immunity level (Moll, 2020). At the end of the pandemic, 60 percent have ever been infected (targeted by the parameterization), and 0.38 percent have died (60 percent infected * 0.64 percent IFR = 0.38 percent). Hence, the final epidemic size as measured by the share of people that ever gets infected is almost twice as large as the minimal level necessary for herd immunity.

Agents' rational adjustment to disease risk can reduce the overshoot in infections. In the figure, the solid blue lines represent the augmented SIRD ('macro-SIRD') model without any containment policy. In response to the pandemic, aggregate consumption (and hours) spontaneously fall by up to 12 percent and track the level of infections, because susceptible agents voluntarily adjust their consumption and labor supply in order to reduce their risk of infection.[6] The voluntary response reduces the pandemic overshoot by about 10 percentage points (about one third) and deaths by about 0.8 percentage points (about one quarter) relative to the standard SIRD.

---

[4] Except for the calibration of the infection probabilities π_{s1}, π_{s2} and π_{s3}, which use steady state consumption and hours worked, the economy is not modeled in the basic SIRD. Hence, there are no predictions for hours and consumption for the basic SIRD in the figure.

[5] The herd immunity level is achieved when the share of remaining susceptible is such that the epidemic's effective reproduction number $R = R_0 \times$ susceptible is less than 1, at which point the epidemic stops. This implies that a share $1 - \frac{1}{R_0}$ must have been infected (ignoring deaths).

[6] Mechanically, the infected also reduce their consumption because their labor becomes 20 percent less productive. This leads only to a small drop in consumption, however: with a peak infection rate of 5 percent, this amounts to a 1 percent reduction in consumption.
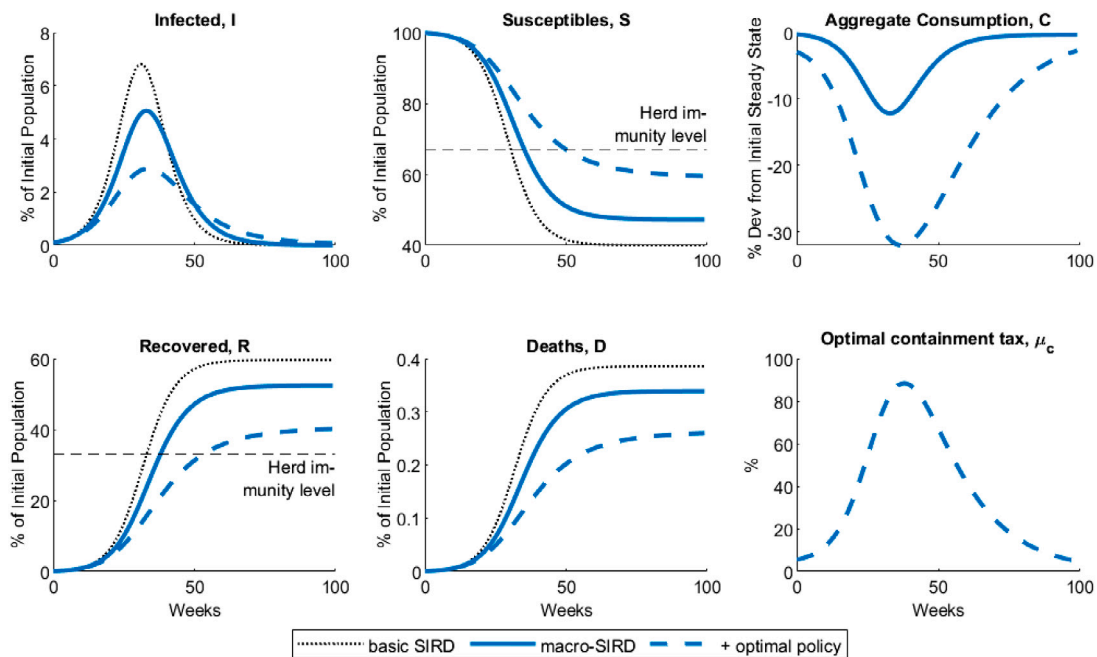
Fig. 2. Benchmark calibration. The figure shows simulated time series for epidemic and economic outcomes for the basic SIRD-model (dotted line), the SIRD-macro model without policy (solid blue line) and the SIRD-macro model with optimal i.e. welfare-maximizing policy (dashed blue line). The dashed horizontal line indicates the minimum share of recovered in the population for the epidemic to end. Parameters in this benchmark calibration are chosen as in Table 1, column 2.

Besides this voluntary adjustment, the social planner can internalize the individual agents' contribution to the pandemic by taxing consumption (μ_{ct} in Eq. (3)) and thereby discouraging economic activity beyond the voluntary adjustment. To choose this 'containment rate', the planner maximizes the sum of the lifetime utility of the agents in the different epidemiological states, weighted by the states' respective sizes subject to the government budget constraint (Eq. (6)). The evolution of the pandemic economy with an optimal – i.e. utility-maximizing – containment policy is represented by the dashed lines in the figure and the optimal containment tax is plotted in the bottom right panel. Containment tracks the path of infections and corresponds to a value-added tax on consumption of up to 88 percent during peak infections or 49 percent on average over the first year. As a result, the economy contracts even further, culminating in a drop of hours and consumption by 35 percent. This further slows infections and reduces deaths by an additional 21 percentage points, almost closing the gap between final epidemic size and the herd immunity level.

3.2. Calibrating the model with heterogeneity in income and IFR, a subsistence constraint and non-homothetic preferences

Having reviewed the basic mechanisms of the model, we now use its extended version, which allows for a non-zero subsistence constraint and an external VSL-target to affect agents' and policy-makers response to the pandemic, to examine how the pandemic evolves in countries that differ from the US benchmark in terms of their income and IFR.

3.3. Heterogeneity in income and the IFR

There are many dimensions along which countries differ and which may affect their response – optimal or actual – to the pandemic. We focus here on two dimensions which, a priori, seem salient for a country's response, and for which data are available: (i) the country's income and (ii) the country-specific risk of dying from COVID-19 when infected.

Specifically, for each country's weekly income $Y_{weekly,j}$, we take World Bank GDP per capita figures at 2010 constant dollars and convert them into weekly numbers. We also allow for (limited) variation in

hours worked by income: Following Charmes (2015), weekly hours are set to 30 for high and upper middle income countries and to 50 for low and lower middle income countries.[7] For country j's IFR, we use the values predicted by Ghisolfi et al. (2020). The authors estimate the COVID mortality rate in a wide range of countries using each country's specific age distribution, the prevalence of comorbidities which have been found to increase risk of death from COVID, and adjusting for the health system's ability to treat respiratory infections. The final estimates of the predicted COVID-19 IFR range from 0.37% in Western Sub-Saharan Africa to 1.45% for Eastern Europe

3.4. Preferences

Once we consider countries with incomes far below the United States, it becomes salient to allow the pandemic response to vary (non-linearly) with a country's income. We do this in two ways: First, we admit non-zero values for the subsistence constraint $\bar{c}$ in the utility function. This implies that the elasticity of substitution between leisure and consumption is no longer constant with income as would be the case for log-preferences in the benchmark simulation. Specifically, for $\bar{c} > 0$, the poorer a country and thus the closer its income to the $\bar{c}$, the less agents will respond to a containment tax.

Second, we follow Viscusi and Masterman (2017) in making a country's value of a statistical life a log-linear function of its income with the common cross-country income elasticity of the VSL denoted by η. Specifically, a country's VSL depends on a base VSL and the ratio between its income and the income of the base country via the following function:

$$VSL_j(\eta) = VSL_{base} \times \left(\frac{Y_j}{Y_{base}}\right)^{\eta}$$

As the base country we use the United States because there are reliable empirical estimates for its VSL. For example, Viscusi and Masterman (2017) report a value of \$9.6million. Here, we use \$$VSL_{base}$ =

---

[7] We obtain the latter value from summing the reported weekly hours of paid and unpaid work for men and women in time-use studies from Sub-Saharan Africa since 2010, as summarized in Charmes (2015).

9.3million, which is similar and equals the VSL implied by the model for the US when $\bar{u} = 0$. Y_j and Y_{base} are given by GDP/capita of country j and the US, respectively. From the expression for the VSL in country j, it can be seen that if $\eta < 1$ countries that are poorer (richer) than the base country have a higher (lower) VSL-to-income ratio than the base case. More generally, countries with different incomes respond differently to the pandemic when the common income elasticity, η, differs from unity because their VSL-income ratios differs.

The VSL introduced in Eq. (2) will in general differ from the $-value of life-time discounted utility over steady-state consumption and leisure and thus requires $\bar{u} \neq 0$. Hence, we need to adjust $\bar{u}$ in the utility function, such that the discounted lifetime utility U in consumption units equals the desired VSL. That is,

$$\frac{U}{u'(c_j)} = \frac{c_j - \bar{c}}{1 - \beta} (\ln(c_j - \bar{c}) - \frac{\theta}{2} n_j^2 + \bar{u}) = VSL_j(\eta),$$

where c_j and n_j are steady-state consumption in the absence of a pandemic in country j. Solving for $\bar{u}$ gives

$$\bar{u} = VSL_j(\eta) \frac{1 - \beta}{c_j - \bar{c}} - (\ln(c_j - \bar{c}) - \frac{\theta_j}{2} n_j^2) = U(1 - \beta) - (\ln(c_j - \bar{c}) - \frac{\theta_j}{2} n_j^2). \quad (7)$$

Thus for the model to deliver a particular $VSL_j(\eta)$, $\bar{u}$ must equal the difference in the per-period value of the discounted life-time utility implied by the desired $VSL_j(\eta)$ and the per-period utility from steady-state consumption and leisure.

We remain agnostic as to the precise values of the subsistence constraint $\bar{c}$ and the income elasticity of the VLS, η. Instead, we first examine how the model's outcomes vary for country's at different income and IFR levels as the preference parameters traverse a grid of plausible values (Section 3.4.1), in order to then structurally estimate $\bar{c}$ and η on this grid (Sections 5.1 and 5.2).

In setting the range for the preference grid, we are motivated by existing estimates for the absolute subsistence constraint and the income elasticity of the VSL. For the subsistence constraint, we consider values from $0 to $600, covering the estimates presented in Duarte and Restuccia (2010; ∼ $0) and Bryan et al. (2014; ∼ $50). The upper bound corresponds to the 5th percentile of yearly GDP/capita across countries where such data is available, which is close to, but below, the World Bank's $2.15\frac{\$}{\text{day}}$ poverty line.[8] To allow for the possibility that poor countries may respond more to a containment tax than rich ones, we also consider negative values of $\bar{c}$, namely $\bar{c} = \$-200$. In this case, the parameter $\bar{c}$ could be interpreted as a kind of basic consumption level that is always achieved (akin to an unconditional safety net or home production) even when households do not purchase consumption.[9] Note that, for simplicity, we will refer to $\bar{c}$ as the subsistence constraint in the rest of the paper.

For the income elasticity of the VSL, we consider values between 0.7 and 1.3. This range encompasses the values for the income elasticity advocated by Viscusi and Masterman (2017), who estimate η to lie between 0.8 to 1.2. For these parameter ranges the utility equivalent of the targeted VSL for poorer countries is larger than the steady-state discounted life-time utility from consumption and leisure in the model, i.e. $\bar{u} > 0$.

With this specification and range for $\bar{c}$ and η, we obtain the special case of constant elasticity between leisure and consumption as well as homothetic preferences between lives and livelihoods when $\bar{c} = 0$ and $\eta = 1$. That is, for this combination of preferences, all countries would set the same percentage value added tax and experience the same percentage drop in consumption. For all other combinations of $\bar{c}$ and η, a country's response to the pandemic is not proportional to income. To be sure, applying the original ERT model with $\bar{u} = 0$ to countries with lower income than the United States also generates non-homotheticity, but not in a desirable way. Specifically, a $\bar{u} = 0$ implies values for the VSL-income ratio in poorer countries that are unreasonably low. For example, in Uganda, the VSL-income ratio would be 52, one third of the US' VSL-income ratio, producing an income-elasticity of 5 for the VSL, far outside of our grid. This would make living is an extreme luxury good.

3.4.1. Comparative statics: how do income and the ifr affect agents' and policymakers' response to the pandemic?

We now present comparative statics to show that the model outcomes vary substantially and systematically with the dimensions of observed cross-country heterogeneity, income, and IFR. We also find that the direction and strength of this variation, especially for income, depends on the direction and strength of non-homotheticity.

To do so, we construct two pairs of representative countries: (i) a 'rich-poor' pair, where pre-pandemic steady-state incomes equal the World Bank's minimum thresholds for high and middle income countries (GDP/capita = $12,696 for the 'rich' and GDP/capita = $1,045 for the 'poor' country) and IFR is held constant and set to the median IFR in our sample of countries (0.5%); and (ii) a 'high-low IFR' pair with the IFRs equal to the 10th and 90th percentile of the global predicted distribution in Ghisolfi et al. (2020) (0.3 percent for the 'low-IFR' and 0.9 percent for the 'high-IFR' country) and pre-pandemic steady-state income equals the median income level in our sample of countries (GDP/capita = $5,853). We set weekly hours to 50 for both pairs of countries.

We examine the effect of the non-homothetic preferences and the subsistence constraint separately by varying one preference parameter at a time. First, we examine the VSL-income ratio channel by varying η over a range from 0.7 to 1.3, while fixing $\bar{c}$ at zero (which implies a constant response to a given tax across countries). Second, we examine the subsistence constraint channel by varying $\bar{c}$ from $-200 to $600, while fixing η at one (which implies a constant VSL-income ratio across countries).

With these inputs, the model's parameters for the representative countries in each of the pairs described above are calibrated as follows: the weekly discount factor is set to $0.96^{1/52}$ as in the benchmark calibration. Labor productivity is set to $A_{repr} = \frac{Y_{weekly,repr}}{H_{weekly,repr}}$, which produces steady state weekly consumption of $c^*_{weekly,repr} = \frac{A}{(H_{weekly,repr})}$. The disutility of labor is found by inverting the first-order condition for hour's worked when there is a subsistence constraint in the utility function, which gives $\theta_{repr} = \frac{A_{repr}}{A_{repr}(H_{weekly,repr})^2(1 - \frac{\bar{c}}{H_{weekly,repr}})^2}$. Given these parameters and values for η and $\bar{c}$ on the grid, $\bar{u}$ is found from Eq. (7).

The time to recovery and the productivity drop from infection are set to the same values as in the benchmark case. The scaling parameters π_{sk} in the probability of infection are calibrated with the representative country's steady state consumption using the same procedure described in Section 3.1. (We summarize the calibration for the pairs of representative countries in column 3 of Table 1).

The comparative statics results for the rich-poor and high-low IFR pairs of countries are shown in Fig. 3. Each column of the figure graphs one of three measures of the economic response (peak drop in hours worked with no containment tax, peak government response in the form of value added tax, and peak drop in hours with the optimal government response over the first year of the pandemic) and

---

[8] We do not include the $2.15\frac{\$}{\text{day}}$ poverty line in our grid since the relatively high (uniform) value for $\bar{c}$ would imply a negative argument to the log-utility function for the poorest countries and thus lead to their exclusion from the simulations.

[9] Negative values of $\bar{c}$ are usually considered for sub-aggregates of consumption, for example, in models of structural transformation where aggregate consumption is disaggregated into agricultural, manufacturing and service consumption, where it is assumed that households always consume a nonnegative amount of services regardless of their consumption decisions (see Herrendorf et al. (2014)). While there may be less empirical support for a negative $\bar{c}$ in the utility function for aggregate consumption, we nevertheless use this modeling device to consider the possibility that the response to a containment tax could, in principle, be decreasing in income.

two measures of the pandemic's evolution (cumulative infections and cumulative deaths from COVID-19 after the first year of the pandemic) on the vertical axis as function of the parameters on the horizontal axis.[10] In the first rows of panel (a) and panel (b), the variable on the x-axis is the income-elasticity of the VSL, η, (with $\bar{c} = 0$ fixed) while in the second rows in panel (a) and panel (b) the variable on the x-axis is the subsistence level, $\bar{c}$ (with $\eta = 1$ fixed). Each subgraph in panel (a) contains two curves representing the two income scenarios, and each subgraph in panel (b) contains two curves representing the two IFR scenarios. The lines' slopes thus represent the partial derivative of the given outcome with respect to the parameter on the horizontal axis (holding income and IFR constant). The height difference of the lines, on the other hand, measure how the model's outcomes vary across rich and poor and low- and high-IFR countries (for given values of the preference parameters η and $\bar{c}$).

Our main concern is to explain the observed gap in pandemic response between rich and poor countries, or high- and low-IFR countries, which may be functions of the unobserved preferences. In short, we are concerned with the relative height of the two lines in each panel of Fig. 3 over the underlying grid of preferences.

First, in the top row of Fig. 3(a), the individual response in absence of a tax, optimal tax, and total response are all greater in poor than in rich countries for $\eta < 1$, smaller for $\eta > 1$ and independent of income for $\eta = 1$. By way of intuition, note that when $\eta < 1$ the VSL-to-income ratio is higher for poorer countries, and thus (in absence of a subsistence constraint which equals zero in the top row) they are more willing to reduce economic activity to avoid pandemic deaths. As a result, infections and deaths are lower in poor countries when $\eta < 1$ (and vice versa if $\eta > 1$).[11]

Second, in the bottom row of Fig. 3(a), the individual response is greater in poor than in rich countries for $\bar{c} < 0$, smaller for $\bar{c} > 0$, and independent of income for $\bar{c} = 0$ (and $\eta = 1$). This happens because with a higher subsistence constraint, the marginal utility cost of reducing consumption or hours worked to avoid deaths is higher — especially in poor countries whose consumption levels are close to the subsistence constraint already in pre-pandemic times.

The optimal policy, however, leans against this individual response, and is stricter in poor countries for a positive subsistence constraint. This happens because the policymaker needs to counteract the response of the infected agents, who—because of the subsistence constraint—increase their labor supply relative to the recovered in order to compensate for the productivity drop that arises as a consequence of infection. Similarly, for a given infection level, the subsistence constraint also mutes the voluntary response of the susceptible.[12] The policymaker is now faced with a world where agents find a reduction in hours and consumption more costly for a given infection level, but where there are more infections relative to a zero subsistence constraint. This leads the policymaker to set a stricter containment policy, albeit not strict enough to completely overturn the weaker response of the susceptible (since the policymaker also cares about the agents' subsistence constraint).

The net effect is a total drop in hours worked that is smaller in poor countries the higher is the positive subsistence level, while the response of rich countries does not vary much with the subsistence constraint. As a result, the model predicts worse epidemic outcomes in poor than rich countries, i.e. higher infections and deaths, when the subsistence constraint is high.

Third, in Fig. 3(b), we see that the individual response in absence of a tax, the optimal tax, and hence the total response are always weaker in low IFR countries than in high IFR countries for all values of the unobserved preference parameters η and $\bar{c}$. As a result, infections, which are more costly the higher the IFR, are decreasing in the IFR. In the end, this is not enough for high IFR countries to reduce deaths below those of low IFR countries though, even in the SIRD-Macro model with optimal policy.

In sum, we have shown that the optimal policy prescriptions and agents' response to it strongly depend on the IFR and income, as well as on the underlying parameters. Thus, both sets of extensions we propose appear important for understanding pandemic policy-making outside of high-income countries.

Specifically, the ranking of rich and poor countries' responses will depend on the income elasticity of the VSL and the subsistence level of consumption. Put another way, poorer countries should set laxer policies because of their lower income only for certain parameter values. And in particular, the intuition that policy should be laxer because of subsistence constraints may fail because it neglects the fact that these constraints themselves raise infections, which calls for stricter policies all else equal. Finally, the comparative statics also show that a lower IFR, regardless of countries' preferences, always leads to laxer policy.

4. Empirical evidence

We now examine the empirical counterparts of the macro-pandemic model outcomes presented in the previous section. Specifically, we show how actual containment policies, economic contractions, infections, and deaths differed between rich and poor as well as high- and low-IFR countries during the first year of the pandemic.[13]

4.1. Data

We use the following outcomes (or proxies) measured throughout 2020. As a measure of national containment policies, we use the lockdown index and the stringency index from the Oxford University Blavatnik School of Government (Hale et al., 2021). The lockdown index captures restrictions on movement and international travel, ranging from 0, when there is no lockdown, to 3, for the most severe lockdown. The stringency index calculates an average of all containment measures a country has taken, including school closures, workplace closures, cancellation of public events, restrictions on gathering size, closure of public transport, and stay-at-home requirements. Varying between 0 and 100, these data are available for 160 countries at weekly frequency. For both indices, we compute the country-level maximum between January and December 2020 to allow for differential timing of the pandemic across countries.

To measure the reduction in economic activity, we use data from the Google Mobility Report (Google LLC, 2021), which tracks how

---

[10] We compare the maxima/cumulative values of such outcomes through time as these are sufficient statistics for the "strength" of the model's response to the pandemic.

[11] The yellow and blue lines in panel (a) of Fig. 3 cross at values slightly different from unity, since even when we scale the VSL linearly with income ($\eta = 1$), differences in the assumed steady-state values of consumption let the epidemiological parameters differ between scenarios. Specifically, since we assume final epidemic size to be constant across countries, but infections to increase with hours worked and consumption expenditures, much lower levels of the latter in LIC imply substantially larger values for π_{s1} there. To address this, we could, in principle, hold epidemiological parameters constant across contexts, and assume a function converting differing hours worked and consumption expenditure into identical levels of actual exposure. Since, however, the intersections' distances to unity are rather small, we stick with the original formulation.

[12] For the susceptible, the effect of a subsistence constraint is ambiguous, the subsistence constraint has two effects: (i) for given infections, it makes their voluntary response more muted, but (ii) the increased activity of the infected encourages a stronger voluntary response. For our parameter values, the first effect outweighs the second, and hence the individual response is decreasing in $\bar{c}$.

[13] We concentrate on the first year of the pandemic before the widespread roll-out of vaccines in high-income countries, which may have provided alternative justification for different policies.

(a) Varying income
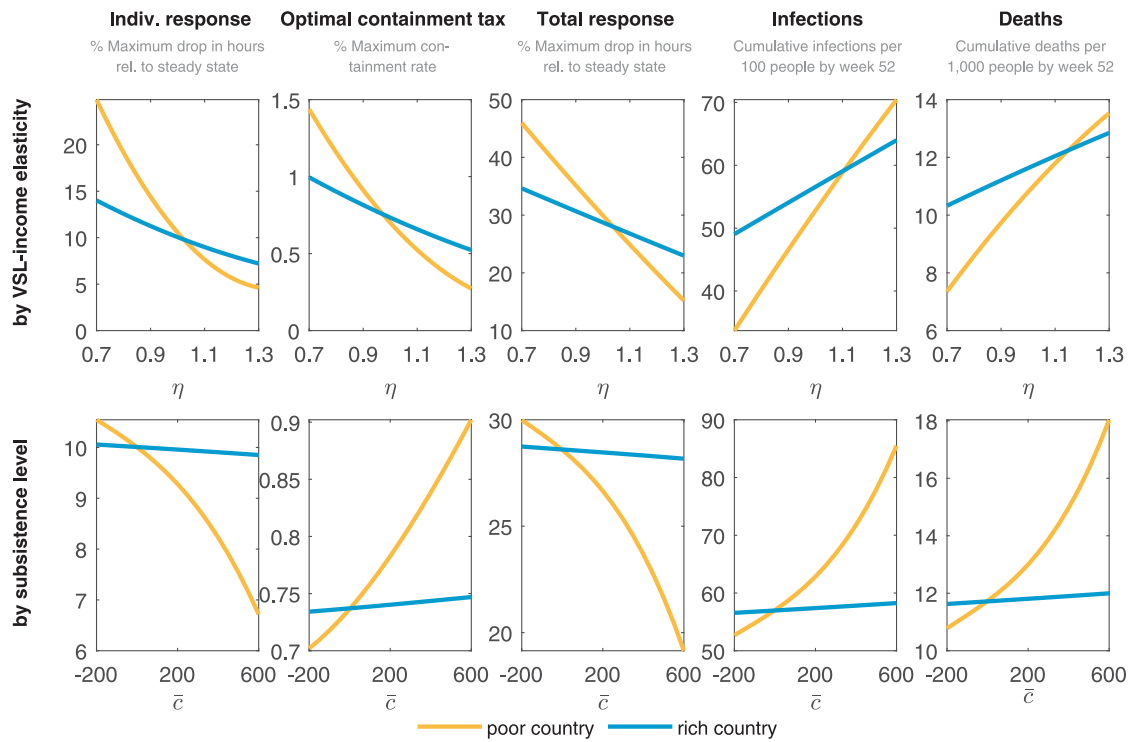
(b) Varying IFR
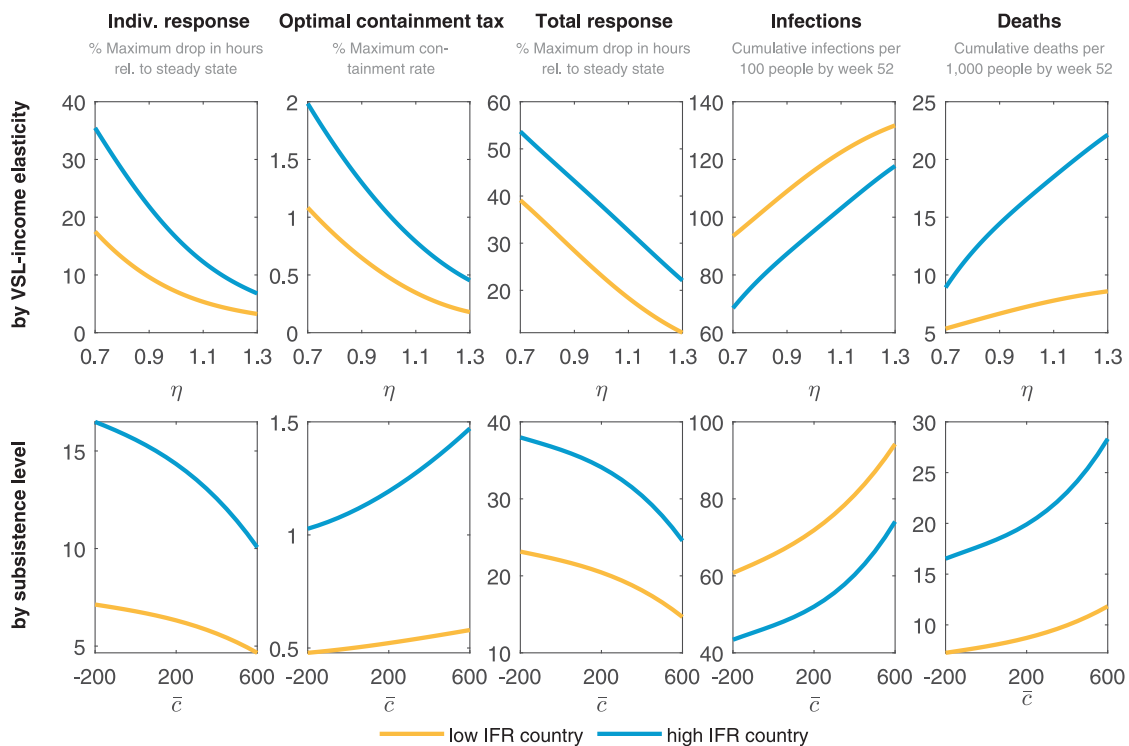
Fig. 3. Comparative statics of economic response by varying $\bar{c}$ and η. Panel (a) shows simulations of five model outcomes (three economic and two epidemic) across two sets of parameter ranges (VSL-income elasticity $\eta \in [0.7, 1.3]$ in top row; subsistence constraint $\bar{c} \in [\$-200, \$600]$ in bottom row), comparing model configurations differing in their income while fixing the IFR at the median across countries. Panel (b) shows the same statistics for the same parameter ranges, comparing model configurations differing in their IFR and holding income constant at the median across countries. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

often GoogleMaps users frequent specific locations relative to pre-pandemic levels. Locations are classified into six categories: grocery and pharmacy; parks; transit stations; retail and recreation; residential; and workplaces. In this analysis, we use workplace attendance as a proxy for hours worked. This data is available at weekly frequency for 125 countries. For each country, we compute the maximum drop in mobility in 2020.[14] As an alternative indicator, we also compute differences in per capita consumption from national accounts data for 2020, available annually for 135 countries from the World Bank World Development Indicators, 2022.

To measure infections, we use the SeroTracker dataset compiled by Arora et al. (2021), which collects COVID-19 seroprevalence studies across a large number of rich and poor countries. The data set includes some studies drawn from convenience samples (such as medical personnel, vendors, or taxi drivers), who may be either more or less likely to be infected than the general population. To avoid this bias, we select studies where the sample of interest was categorized as either "Blood donors" or "Household and community samples". We then compute a weekly average seroprevalence by country. The result is a dataset of 195 country-week observations for 28 countries.[15]

To measure deaths, we would ideally like to have data on cumulative COVID-19 mortality in the first year of the pandemic. While official death tallies exist for almost all countries on a weekly basis (see e.g., Hale et al. (2021)), there is strong evidence they severely undercount actual COVID-19 deaths, especially in poorer countries (The Economist, 2022). Thus we follow the common practice of relying on excess all-cause mortality estimates, drawing on the estimates for cumulative excess mortality compiled for 191 countries by Wang et al. (2022) and, specifically, on their estimate of the country-specific ratio of excess mortality to official mortality at the end of 2021. The remaining challenge we face is that Wang et al. (2022) report cumulative measures for the end of 2021, whereas our focus is on the first year of the pandemic. To arrive at cumulative excess deaths at the end of 2020, we multiply the cumulative officially recorded COVID-19 deaths at the end of 2020 by the ratio computed by Wang et al. (2022).[16] The result is an estimate of cumulative excess mortality at the end of the first year of the pandemic for 174 countries.

We classify countries as high- or low-income using World Bank GDP/capita figures at 2010 constant dollars. A country is classified as high-income if its GDP is above the median in the data and as low-income otherwise. Furthermore, we classify countries as high- or low-IFR on the basis of their predicted IFR according to Ghisolfi et al. (2020). A country is classified as high-IFR if its predicted IFR is above the median in the data and as low-IFR otherwise.

---

[14] It is possible that by focusing on smartphone users, the index may not represent the average response for the overall population. However, given findings that poor citizens are less able to adjust labor supply in response to shocks (Desai and Pramanik, 2020; Chetty et al., 2021), and assuming that poor individuals also have lower rates of smartphone ownership in low-income countries, the effect we find should be a lower bound of the difference in response between high- and low-income countries.

[15] The SeroTracker aggregates published peer reviewed research articles, preprints, reports, and media (unpublished grey literature). In our dataset, 74 (37%) observations come from peer-reviewed, published studies (at the time of writing), 64(32%) preprints, 40 (20%) official reports, and 20 (10%) come from news or media. Rich and poor countries have the same number of data points coming from published articles or media, but no poor country has data points coming from official reports, while they have double the number of data points coming from preprints as rich countries. We also confirm that the timing of the seroprevalence studies does not differ between high and low income countries. The sign and magnitudes of the estimates do not vary if we include 1. only the seroprevalence data coming from peer-reviewed articles 2. peer-reviewed or preprint, 3. peer-review, preprint or official reports.

[16] As the ratio of excess to official mortality in rich countries can also take values below one (and, in some cases, can even be negative), we bound the minimum number of estimated deaths to be at least as high as the number of recorded deaths. The results are very similar without this adjustment.

4.2. Empirical results

We now estimate the gap in outcomes between high and low-IFR countries and rich and poor countries by regressing each outcome $Y = \{$Containment policy (P), Drop in Economic Activity (ED), Infections (I), Deaths $(D)\}$ on an indicator variable $\mathbb{I}_{X>\text{median}}$ for whether a country falls above or below the global median for either the IFR $(X = \text{IFR})$ or for income per capita $(X = \text{GDP/capita})$[17]:

$$Y_i = \alpha + \beta\mathbb{I}_{X>\text{median}} + \varepsilon_i \qquad (8)$$

For the regressions using peak containment, peak drop in economic activity, and cumulative excess deaths, data are at the country-level with one observation per country. For seroprevalence, the data comprise an unbalanced panel, with one observation per country-week when a seroprevalence study was reported. In this regression, standard errors are clustered at the country level.

The results are summarized in Table 2. In columns (1) and (2), we report the average outcomes in high- and low-IFR countries, respectively. Column (3) reports the difference (expressed as percentage of the high-IFR countries' outcome) and in column (4), the p-value of the difference. Columns (5) to (8) report the same statistics for rich and poor countries. Column (9) reports the number of observations for each outcome (see Table 2).

High-IFR countries set on average stricter containment policies than low-IFR countries: maximum lockdown severity was 2.11 (out of 3) and maximum stringency was 84 percent in high-IFR countries, while these numbers stood at the 2.00 and 83 percent in low-IFR countries. Consistent with this, the economies of low-IFR countries contracted less during the pandemic: at its lowest point, mobility had dropped by 56 percent in high-IFR countries and by 49 percent in low-IFR countries. Similarly, annual per capita consumption decreased by 5 percentage points in high-IFR countries and by 2.2 percentage points in low-IFR countries.

On the health side, countries with a high IFR experienced fewer infections, but more deaths than low-IFR countries. The average seroprevalence was 9 percent in high-IFR countries and 17 percent in low-IFR countries, while cumulative excess deaths stood at 96.3 per 100,000 in high-IFR countries and at 49.8 per 100,000 in low-IFR countries.

Comparing rich and poor countries, we find that the latter set the same (if not stricter) containment policies, confirming the graphical analysis in the introduction. Maximum lockdown severity was 2.06 and maximum average stringency was 84 percent in both sets of countries. Despite similar containment, the economy contracted 16 percent less in poor countries, and the average consumption drop was 4 percentage point lower.

The pandemic's death toll was higher in richer countries during its first year (this changes in 2021 due to the introduction of the vaccines), though the difference between rich and poor countries is not as pronounced as the difference between high- and low-IFR countries. Cumulative excess deaths reached 77.9 per 100,000 in rich countries and 68.1 per 100,000 in poor countries. Infection rates, on the other hand, were higher in poorer countries, with a seroprevalence of 20 percent compared to 9 percent.

---

[17] Note that we discuss here the unconditional differences, even if we acknowledge that IFR and GDP are not independent. The reason is that the literature and policy debate has focused on the difference between high and low-income countries unconditionally, therefore we choose to target these moments rather than conditional ones when we estimate the model on the full sample of countries. We have also conducted the model estimation on the conditional moments, and our results remain the same.

Table 2
Estimated outcome gaps.

	(1) High IFR	(2) Low IFR	(3) % Gap	(4) P-value	(5) High GDP	(6) Low GDP	(7) % Gap	(8) P-value	(9) Obs
Peak policy difference									
Stringency	84.43	83.37	1	0.59	83.38	84.42	−1	0.50	160
Lockdown	2.11	2.00	5	0.36	2.06	2.06	−0.2	0.90	160
Peak economic activity difference									
Max mobility drop	56.38	48.96	13	0.00	57.48	48.02	16	0.00	126
Avg consumption drop	5.00	2.23	55	0.01	5.53	1.87	66	0.00	135
Cumulative infections difference									
Seroprevalence	0.09	0.17	−82	0.01	0.09	0.20	−131	0.00	195
Cumulative deaths difference									
Excess death rate	96.30	49.74	48	0.00	77.90	68.14	13	0.40	174

Columns (1) and (2) report the average outcomes in high- and low-IFR countries, respectively; column (3) reports the difference (expressed as percentage of the high-IFR countries' outcome); and column (4) reports the *p*-value of the difference. Columns (5) to (8) report the same statistics for rich and poor countries. Column (9) reports the number of observations for each outcome.

5. Matching the model to the data

We now ask whether the macro-pandemic model can generate the pattern of outcomes we see in the data for "sensible values" of the underlying parameters η and $\bar{c}$. The stylized example and comparative statics in Section 3 suggest that it could, but those simulations do not take into account the strong positive correlation between a country's income and its predicted IFR that we observe in the data. Hence we now simulate the model's outcomes conditional on the full joint distribution of income and IFR, i.e. for all countries in our data set.

With this simulated data we can answer the question at hand using indirect inference. That is, we find the set of preference parameters, if any, for which the model produces rich-poor and high-low IFR gaps which match the empirical ones or, more precisely, their signs. We content ourselves with matching signs – rather than magnitudes – of the target moments, because model outcomes are not directly observed in the real world and we rely on proxies (stringency for containment tax, mobility for consumption and hours worked) or variables measured with error (as in the case of seroprevalence and deaths).

5.1. Simulated endogenous model outcomes

To simulate the model-implied rich-poor and high-low IFR gaps, we return to the country-aggregate data presented in the previous section, focusing on the 125 countries for which we have full data for containment, mobility, and deaths over the course of the pandemic.

Specifically, we calibrate the model's parameters for each of 125 countries using the country's predicted IFR, pre-pandemic observed GDP per capita and hours worked and letting η and $\bar{c}$ take on values on a 7×5 where η ranges from 0.7 to 1.3 in steps of 0.1 and $\bar{c}$ ranges from −200 to 600 in steps of 200.[18] The calibration is analogous to what is described in Section 3.4.1, but uses each individual (rather than the representative) country's GDP, hours and IFR. Full details of the calibration are found in column 4 of Table 1.

For each of the 35 combinations of η and $\bar{c}$ on the grid, the model is then solved for each of the 125 countries. Specifically, we create simulated sequences for weekly containment policy (the planner's optimal value added tax), the drop in economic activity (the model's percentage deviation of aggregate consumption from the steady state), infections, and deaths for the first 100 weeks of the pandemic, for a total of 35 sets of simulated data for 125 countries.

For each of these 35 sets of simulated data, we then calculate the size of the gap for peak containment policy, peak economic contraction, cumulative infections, and cumulative deaths between rich
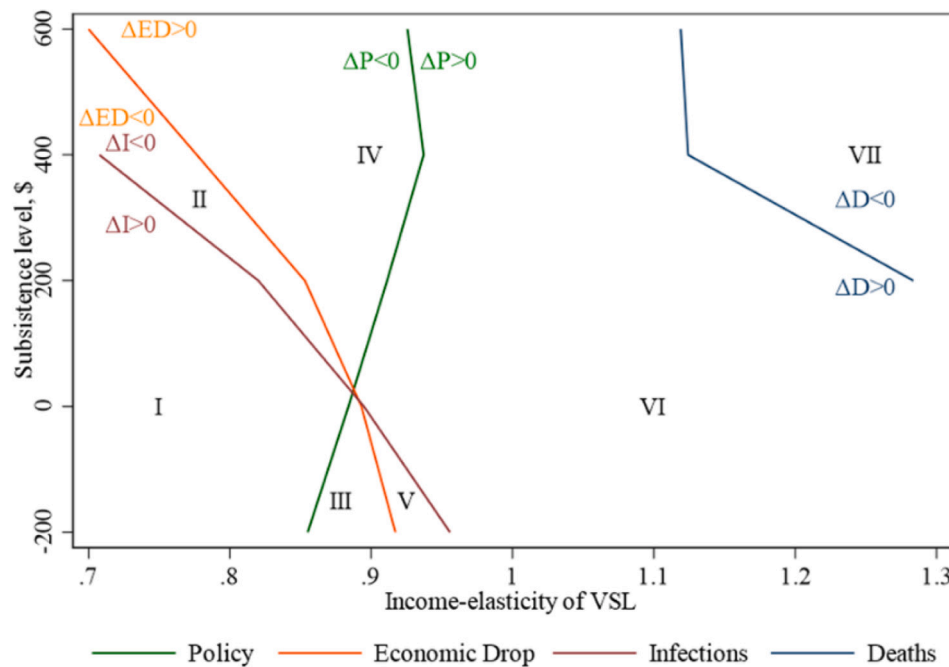
and poor countries and high- and low-IFR countries over the first year of the pandemic, just as in Section 4.2. That is, we calculate for each $Y = \{\text{Containment policy } (P), \text{Drop in Economic Activity } (ED), \text{Infections } (I), \text{Deaths } (D)\}$ and for each $X = \{IFR, GDP\}$:

$$\Delta_X Y(\eta, \bar{c}) = E(Y_i(\eta, \bar{c})|X_i \rangle \text{median}) - E(Y_i(\eta, \bar{c})|X_i \leq \text{median}) \quad (9)$$

In Fig. 4, Panel A, we graph the combinations of η and $\bar{c}$ which produce gaps equal to zero between rich and poor countries in the four outcomes, and indicate regions where the gaps have positive/negative sign. The zero-level curves for the IFR gaps lie outside the parameter space and are therefore not shown. We focus on the zero-lines both because we are only interested in matching the signs of the moments and because the contour lines related to different levels of the gap are effectively parallel (the full contour map for non-zero levels of the gaps is shown in Fig. A.1).

Relative to the comparative statics in Section 3.4.1, the qualitative model predictions remain the same when simulating the model's endogenous outcomes on the full two-dimensional parameter space and for the joint distribution of income and IFR in data. For example, it is still the case that policy is always stricter, and the total drop in economic activity always higher and thus total infections lower in high-IFR countries. Despite this, high-IFR countries always experience more deaths. Together, this results in $\text{sgn}(\Delta P, \Delta ED, \Delta I, \Delta D) = \{+, +, -, +\}$ for the IFR gaps for all values of the preference parameters $\eta, \bar{c}$ (as seen in the full contour map in the Fig. A.1). Similarly, as can be seen from Fig. 4, it remains the case, that rich-poor gaps in policy, economic contraction, infections and deaths can take on positive, zero, and negative values depending on the parameters. And just as before, the rich-poor policy gap is increasing in η and decreasing in $\bar{c}$, (indicated by the upward sloping null line), while the rich-poor gap in the economic contraction, infections and deaths is increasing in both η and $\bar{c}$ (indicated by the downward sloping null lines).

Simulating the model on the two-dimensional parameter space does, however, clarify the trade-offs between the preference parameters in generating the model's outcomes. For example, both an income elasticity of the VSL of $\eta = 0.95$ and a subsistence level of $\bar{c} = 600$, and an $\eta = 0.88$ and a $\bar{c} = 0$, will produce a zero policy-gap between rich and poor countries. In the first scenario, poor countries have a slightly higher VSL-to-income ratio than rich countries, but – because of the high subsistence constraint – a much lower individual response and sensitivity to the containment tax, which the poor country policymaker needs to counteract by setting a stricter policy. In the second case, there is no subsistence constraint and so policymakers do not need to compensate for limited shielding by the susceptible. Their VSL-to-income ratio is much higher than in the first scenario, though, which pushes them towards stricter policies. These two effects cancel each other out to produce the same (zero) policy gap in the two scenarios.

---

[18] Pre-pandemic hours worked are set to 30 for countries above median income and to 50 for countries below, see Section 3.3.

Panel B: Sign patterns

Region	IFR				GDP			
	Policy	Econ. Drop	Infections	Deaths	Policy	Econ. Drop	Infections	Deaths
I	+	+	-	+	-	-	+	+
II	+	+	-	+	-	-	-	+
III	+	+	-	+	+	-	+	+
IV	+	+	-	+	-	+	-	+
V	+	+	-	+	+	+	+	+
VI	+	+	-	+	+	+	-	+
VII	+	+	-	+	+	+	-	-

Fig. 4. Set-identification of the parameters. Panel A combines the zero-level curves for the four rich-poor gaps for policy (ΔP), drop in economic activity (ΔED), infections (ΔI) and deaths (ΔD) (see the top row of Fig. A.1). The zero-level curves for the IFR gaps lie outside the parameter space and are hence not shown. Each zero-level curve partitions the parameter space into two regions, producing either a positive or a negative gap for the respective moment (indicated by $\Delta Y \gtrless 0$ next to the level curve). Together, the four zero-level curves partition the parameter space into seven regions, denoted I to VII. Each parameter region generates a particular sign pattern for the high-low IFR and rich-poor gaps. In Panel B, we display the sign pattern for the high-low IFR gaps and summarize the rich-poor sign patterns shown in Panel A. On the left we show the sign patterns for high-low IFR gaps for policy (column 2), drop in economic activity (column 3), infections (column 4) and deaths (column 5). On the right, we show the sign patterns for the rich-poor gaps for the same four pandemic outcomes.

The slope of the null-lines is also informative about the sensitivity of a rich-poor gap to changes in the unobserved preferences. For instance, the zero-level curve for economic policy shows that rich countries enact stricter containment policies for values of η higher than 0.9, almost independent of $\bar{c}$. In contrast, the zero-level curve for excess deaths depends on both parameters: rich countries experience lower deaths than poor countries in the northeast quadrant, corresponding to high values of η and $\bar{c}$.

Finally, taking into account the correlation between IFR and income in the data does change the model outcomes quantitatively relative to the comparative statics in Section 3.4.1. For example, the zero-line for policy (and economic contraction and infections) no longer passes through $\eta = 1$ (and $\bar{c} = 0$), but is located to the left of the unit elasticity. This happens because income and IFR are highly positively correlated in the global data: being in the bottom half of the income distribution across 174 countries implies a 42 percent higher chance of having an IFR below the median. As a result, at $\eta = 1$ poorer countries have the same VSL-to-income ratio as richer countries, but their lower IFR drives them to set a laxer policy. To be pushed to equality thus requires an $\eta < 1$ (between 0.85 and 0.95 depending on the subsistence level), because this increases the VSL-to-income ratio in the poor country relative to the

rich. For the same reason, the zero-line for deaths is located far to the right of unity, because even when richer countries set stricter policies than poorer countries, they will still experience relatively more deaths because their IFR is on average higher.

5.2. Estimating preferences

Broadly, it is clear from Fig. 4 that the signs of the gaps in policy can be used to (set) identify the underlying preferences and test the model's validity.

Regarding the former, variation in preferences creates systematic variation in the signs of the rich-poor gaps. By implication, these gaps can therefore be used to identify preferences. To be precise, the four zero-level curves in Fig. 4 partition the parameter space into seven connected regions, denoted I to VII, one for each sign pattern that the model is able to generate (as summarized in Panel B of the figure). For example, all parameter combinations in region IV generate $\text{sgn}(\Delta P, \Delta ED, \Delta I, \Delta D) = \{+,+,-,+\}$ for the high- vs low-IFR gaps and $\text{sgn}(\Delta P, \Delta ED, \Delta I, \Delta D) = \{-,+,-,+\}$ for the rich-poor gap. Thus, if we observed this sign pattern in the data, we would conclude that the unobserved preferences lie in region IV. The preferences are thus
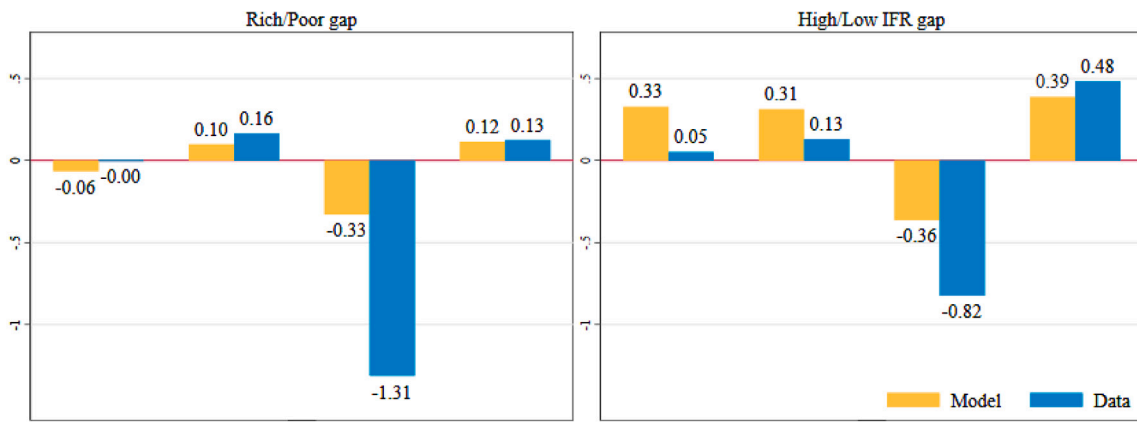
Fig. 5. Alignment of rich-poor and high-low IFR gaps in model and data. The left panel of the figure shows the relative difference between rich and poor countries' $((\bar{Y}_{rich}-\bar{Y}_{poor})/\bar{Y}_{rich})$ modeled (blue) and observed (red) values for four outcomes over the first year of the pandemic. The right panel shows the relative difference between high and low-IFR countries' $((\bar{Y}_{highIFR}-\bar{Y}_{lowIFR})/\bar{Y}_{highIFR})$ modeled (blue) and observed (red) values for four outcomes over the first year of the pandemic. A value close to one implies that the values between poor and rich (low/high IFR) countries are approximately equal and a ratio larger than one implies that the outcome is larger in the poor (low-IFR) country than in the rich (high-IFR country). A difference between blue and red bars implies that the ratios are different between modeled and observed data. 'Peak policy' refers to the peak containment rate in the model and the peak stringency index from Hale et al. (2021). 'Peak economic drop' refers to the reductions in hours worked in the model and the reduction in attendance of workplaces as measured by Google LLC (2021). 'Average Seroprevalence' refers to cumulative infections in the model and estimates compiled in Arora et al. (2021). 'Cumulative deaths' refers to the same in the model and estimates compiled by Wang et al. (2022). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Parameter region estimated to target signs of IFR gaps and rich-poor gaps.

	Data	Model
Estimates		
Parameter region		IV
Income elasticity of VSL, η:		(0.7, 0.95)
Subsistence level, $\bar{c}$:		(50, 600)
Target moments		
High-low IFR gaps		
Policy response	+	+
Drop in economic activity	+	+
Infections	−	−
Deaths	+	+
Rich-poor gaps		
Policy response	−	−
Drop in economic activity	+	+
Infections	−	−
Deaths	+	+

The top panel of the table reports the estimated parameter region and ranges for η and $\bar{c}$ that arise from maximizing the number of sign agreements between the rich-poor and high-low IFR gaps in data and model. The bottom panel reports the signs of the gaps in data and model in the estimated parameter region IV.

set-identified from the rich-poor gaps and the quality of identification depends on precisely what sign pattern one observes in the data (i.e. some regions in the figure are larger than others).

Regarding the latter (model validation), the fact that the model can generate only 7 sign patterns of a possible $2^8 = 256$, and only one sign pattern for the IFR gaps in particular, makes the model falsifiable. Thus matching the signs of the simulated and data moments is not trivial (even for those moments which vary with the unobserved preferences) and, intuitively, the fewer signs we are able to match, the more skeptical we should be that the model generated the data. Specifically, we should reject the model as a data generating process if the sign pattern for the IFR and rich-poor gaps in the data does not correspond to one of the 7 patterns in Panel B of Fig. 4.

Given this heuristic discussion, our estimation boils down to choosing regions for η and $\bar{c}$ to maximize the number of sign agreements between the rich-poor and high-low -IFR gaps in the simulated and observed data. The results of this estimation are shown in Table 3. To recap, in the data, individuals living in poor countries reduced their economic activity less than in rich countries, despite equally strict (or stricter) containment policies. This led to higher infections but a lower

death rate. Countries with lower IFRs set laxer policies, experienced less of an economic contraction, and had higher infections, and lower death rates. The simulated model matches this pattern exactly for parameters in region IV and the number of sign agreements between model and data is thus maximized for any parameter combination in IV. Hence, we estimate lower and upper bounds for η of 0.75 and 0.9 and for $\bar{c}$ of \$50 and \$600. The fact that we match the sign pattern in the data exactly indicates a good fit between model and data.

5.3. Comparing estimated model and data

Having used the congruence of signs to estimate the unknown preferences, we now examine how well the structurally estimated model matches the actual magnitudes of the differences in the core outcomes across rich and poor and high- and low-IFR countries in the data. Within the range of parameter values consistent with the *sign* of our empirical estimates, we focus here on plausible values of both the VSL elasticity and the subsistence constraint, namely $\eta = 0.9$ and $\bar{c} = \$400$, which yield *magnitudes* for policy stringency, the contraction in economic activity, and cumulative death rates that are close to what we observe in the data. For each of the four model-generated outcomes, we calculate the relative rich-poor gaps $((\bar{Y}_{rich}-\bar{Y}_{poor})/\bar{Y}_{rich})$ and relative high/low-IFR gaps $((\bar{Y}_{highIFR}-\bar{Y}_{lowIFR})/\bar{Y}_{highIFR})$ and compare each gap to its data counterpart.

Fig. 5 shows that the magnitudes of the relative rich-poor gaps line up closely between model and data, with the exception of the infection data, where infection rates in poor countries were much higher than our model predicts. If anything, this discrepancy makes the policy response of low-income countries appear even more reasonable.

The magnitudes of the relative high-low IFR gaps diverge somewhat more between model and data. In particular, the model predicts a larger policy gap than what we see in reality. This is perhaps not surprising since our model assumes that policymakers know the relevant IFR, thus abstracting from the fact that there was considerable uncertainty about this parameter, especially at the beginning of the pandemic. Moreover, the pandemic started in countries with relatively high IFRs, which may have influenced policy-making in countries with lower IFRs. Including these considerations in the model would lead it to predict more similar policies in high- and low-IFR countries.

Overall, the broad congruence between model predictions and data for the rich-poor and the high-low IFR gaps across the four outcomes

makes us confident that even a relatively parsimonious model such as ours can be used to rationalize pandemic policy-making and make inference about policy-makers' preferences.

6. Discussion

Do the structurally estimated preferences indicate that poor countries reacted too strongly to the pandemic because they put too high a value on life over livelihoods? If our results implied values of η or $\bar{c}$ that were inconsistent with what we know from the existing literature, for example a VSL that is independent of income or a negative subsistence constraint, we would be inclined to question the wisdom of pandemic policy-making in low-income countries. The range we estimate for the preferences suggests there is no reason to do so. On the contrary, the elasticity we estimate implies a VSL that increases strongly but less than proportionally with income. It also lies squarely in the range estimated in the academic literature and adopted by policymakers. For example, Viscusi and Masterman (2017) collect VSL estimates from contexts varying by income and estimate the cross-country income elasticity of the VSL to lie between 0.8 and 1.2. A review by the OECD (2012) recommends the use of an income elasticity of 0.8.

The estimated elasticity is also consistent with other observed policy choices. For example, the amounts governments and citizens spend to finance their health system relative to income suggest that poorer countries are indeed willing to put relatively more money into saving lives than increasing livelihoods than richer countries. In our sample, we find an income-elasticity of health expenditure of 0.85 among countries categorized as low- and lower-middle income by the World Bank, while the elasticity in the whole sample is around 1.1. More detailed analyses – for example, Farag et al. (2012) – find elasticities lower than unity using either cross-section or time-series variation, which qualifies health expenditure as a non-luxury good.

Similarly, the existence of a subsistence level and thus an elasticity of substitution between consumption and leisure that is increasing in income, is widely accepted. Our finding of $\bar{c} > \$50$ per person per annum is consistent with this. If one wants to relate $\bar{c}$ to an absolute poverty line below which consumption must not fall, then the upper bound of $600 corresponds to about $1.6 per day. Although subsistence levels and poverty lines are certainly not invariant across contexts, such a value seems well within the plausible range of a bare minimum consumption level, below which agents might be unwilling or even unable to trade off consumption to reduce COVID-19 exposure.

7. Conclusion

In early 2020, many countries enacted strict containment policies, imposing economic pain to reduce the spread of disease. Especially during the first stages of the pandemic, the stringency of these policies was quite similar in rich and poor countries.

In this paper, we attempt to make sense of these policy choices from the point of view of national policymakers maximizing social welfare, defined as an aggregate of individual utilities. To do so we follow the emerging COVID-19 economics literature by embedding a standard compartmental model from epidemiology (with agents moving between four categories: susceptible, infected, recovered, or deceased) within a simple macroeconomic framework. On the economic side, agents make decisions about consumption and labor supply, each of which exposes them to an increased risk of infection and, hence, death, but fail to internalize the externality their choices have on others. Policymakers can tax consumption to discourage risk-taking and achieve the social optimum.

We focus on two ways that rich and poor countries differ which we hypothesize are salient for understanding their respective pandemic policy responses: their incomes, of course, and their infection fatality rate (IFR). With their younger populations, poorer countries faced lower IFRs, which in all versions of our model point unambiguously

toward laxer containment policies. We expand on earlier work to allow income to affect agents' and policymakers' choices in two ways. First, as explored in the large literature on the value of a statistical life, poorer countries may place a higher (or lower) monetary value on life in proportion to their income. This is mostly an empirical question: in our theoretical framework, a higher VSL as a share of per capita income points unambiguously toward stricter containment policies, tilting the trade-off between lives and livelihoods toward the former. Second, we also allow for subsistence constraints, such that individuals below some consumption threshold become qualitatively less willing to forego consumption to reduce disease risk. As we show, the impact of subsistence constraints on optimal behavior and policy is more nuanced.

Empirically, we estimate these underlying preference parameters using measures of policy severity (proxied by an index of stringency, aggregating various policy levers such as school closures, stay-at-home orders, etc.), agents' economic behavior (proxied by Google's mobility index), COVID-19 infections from seroprevalence studies, and measures of excess mortality during the pandemic. We see that while poorer countries enacted equal or stricter containment policies, they also witnessed less acute drops in economic activity. Dividing countries by their predicted infection fatality rate, we see (as the model unambiguously predicts) less strict policies and smaller drops in economic activity in countries with lower IFRs, as well as higher infection rates combined with lower mortality rates.

Calibrating the model for 125 countries and a grid of possible VSL and subsistence constraint parameters, we show that the observed empirical patterns are consistent with the model when two conditions hold: (a) the income-elasticity of the VSL is less than one, such that poorer countries have a proportionally higher value on life and (b) the subsistence constraint is strictly greater than zero. Both of these features appear consistent with empirical literature from various sources prior to the pandemic.

It is important to note that we are not in a position to make any claims about whether a given policy in a given country was or was not optimal from a welfare-maximizing perspective. Rather, we find that a model that balances economic and disease considerations using an income specific value of a statistical life and a subsistence constraints appears consistent with the empirical data. Within such a model, we identify the forces that might push poorer countries toward relatively stricter policies, despite their lower infection fatality risk. In combination, these features can explain why a poor country might impose equally strict or stricter containment measures as a rich country, despite that — and in fact, because — agents in the former face subsistence constraints, making it harder for them to comply with containment.

Data availability

Replication packages for the tables and figures in the article are available https://www.dropbox.com/sh/pklj09epvcannf1/AAD9Wi3vf 1CK30F0nD71jJDpa?dl=0.

Appendix

In Fig. A.1, we plot the selected simulated moments on the two-dimensional parameter grid by way of a contour map. In the top row, the contour lines in each panel show the combinations of parameters that produce a given rich-poor gap. In the bottom row of Fig. A.1, we repeat the exercise for the gap in outcomes between high- and low-IFR countries. Of particular importance are the zero-level curves (indicated by a black line) in each panel, which allow us to partition the parameter space into two regions: one region containing parameter combinations that induce a positive gap in the outcome (shaded red with darker colors indicating larger positive values), and one region where the parameter combinations induce a negative gap (shaded blue with darker colors indicating larger negative values) (see Fig. 4, panel B, Table A.1).
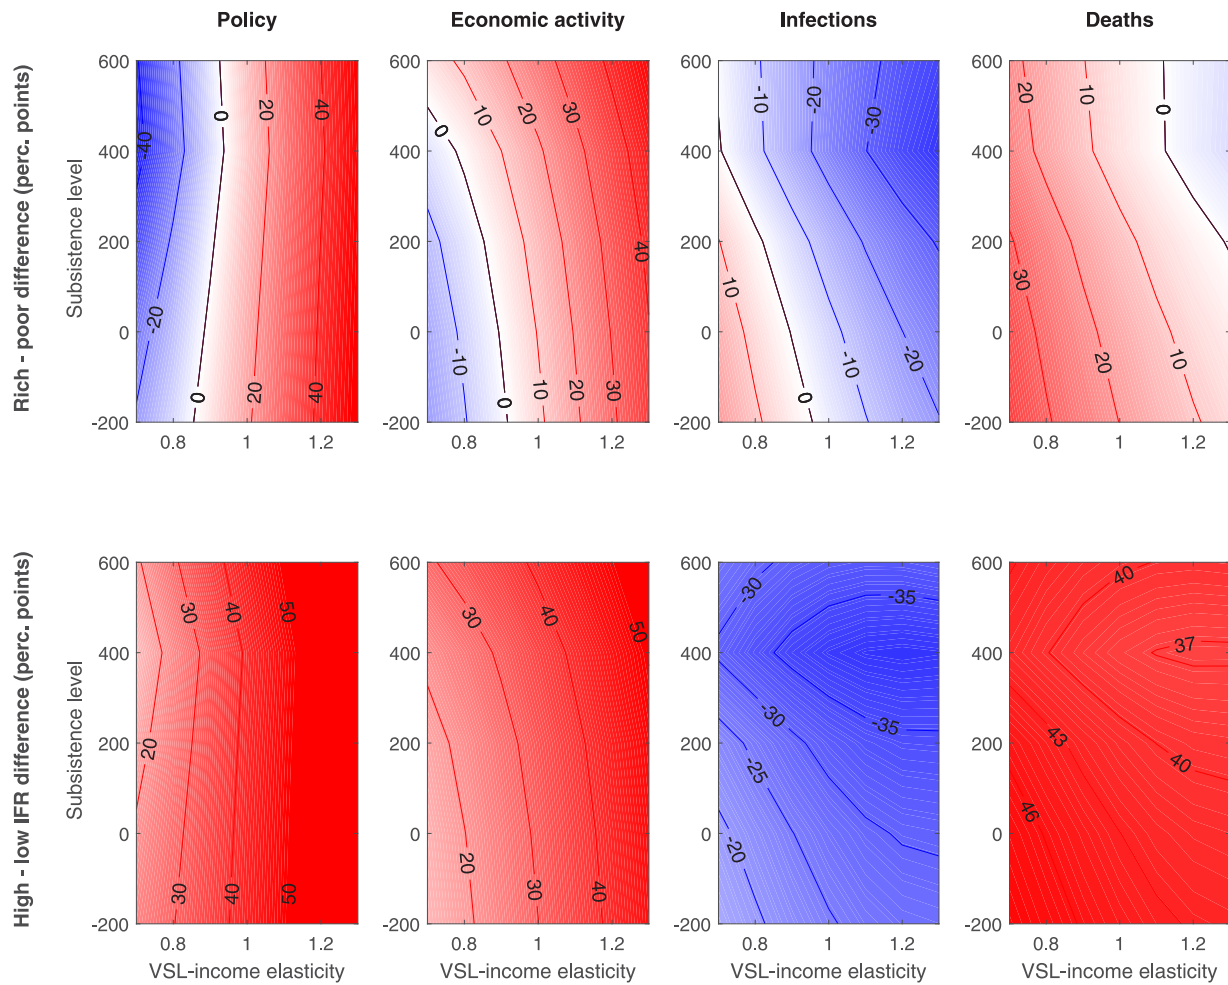
Fig. A.1. Contour maps of outcome differences by rich vs. poor and high vs. low IFR countries. The top row of the figure shows percentage-point differences in model outcomes between rich and poor countries across a parameter grid including the subsistence level $\bar{c} \in [\$-200, \$600]$ on the vertical axis and the VSL-income elasticity $\eta \in [0.7, 1.3]$ on the horizontal axis. 'Policy' refers to the peak policy during the first year of the pandemic. 'Economic activity' refers to peak contraction in hours worked. 'Infections' and 'Deaths' refer to cumulative infections and deaths over the first year of the pandemic, respectively. We classify rich and poor countries by whether they are above or below the global median GDP per capita across countries from the World Bank World Development Indicators. For parameter combinations falling into blue-shaded areas, our model predicts a higher value (e.g. a stricter peak policy, more infections) for rich countries, and vice-versa for red-shaded areas. The bottom row repeats the exercise for the gap in outcomes between high- and low-IFR countries, where we again classify countries by their IFR relative to the global median using data from Ghisolfi et al. (2020). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table A.1
Countries included in analyses.

Country	Stringency	Mobility	Income	Seroprevalence	Deaths	Model	Country	Stringency	Mobility	Income	Seroprevalence	Deaths	Model
Afghanistan	X	X			X	X	Cameroon	X	X	X		X	X
Albania	X		X		X		Canada	X	X	X	X	X	X
Algeria	X		X		X		Central African Rep.	X		X		X	
Angola	X	X	X		X	X	Chad	X		X		X	
Antigua & Barbuda		X			X		Chile	X	X	X		X	X
Argentina	X	X	X		X	X	China	X		X	X	X	
Armenia			X		X		Colombia	X	X	X	X	X	X
Australia	X	X	X	X	X	X	Comoros			X		X	
Austria	X	X	X	X	X	X	Congo Rep.	X		X	X	X	
Azerbaijan	X				X		Costa Rica	X	X	X		X	X
Bahrain	X	X			X	X	Cote d'Ivoire	X	X			X	X
Bangladesh	X	X	X		X	X	Croatia	X	X	X		X	X
Barbados	X	X			X	X	Cuba	X		X		X	
Belarus	X	X	X		X	X	Cyprus	X		X		X	X
Belgium	X	X	X	X	X	X	Czech Rep.	X	X	X		X	X
Belize	X	X	X		X	X	Congo DRC	X		X		X	

(continued on next page)

Table A.1 (*continued*).

Country							Country						
Benin	X	X	X		X	X	Denmark	X	X	X	X	X	X
Bhutan	X		X		X		Dominican Rep.	X	X	X	X	X	X
Bolivia	X	X	X		X	X	Ecuador	X	X	X	X	X	X
Bosnia & Herzegovina	X	X	X		X	X	Egypt	X	X	X		X	X
Botswana	X	X	X		X	X	El Salvador	X	X	X		X	X
Brazil	X	X	X	X	X	X	Equatorial Guinea			X		X	
Brunei	X		X		X		Estonia	X	X	X	X	X	X
Bulgaria	X	X	X		X	X	Eswatini	X		X		X	
Burkina Faso	X	X			X	X	Ethiopia	X		X	X	X	
Cabo Verde	X	X	X		X	X	Fiji	X	X			X	X
Cambodia	X	X	X		X	X	Finland	X	X	X		X	X
France	X	X	X	X	X	X	Liberia	X				X	
Gabon	X	X	X		X	X	Lithuania	X	X	X		X	X
Georgia	X	X	X	X	X	X	Luxembourg	X	X	X	X	X	X
Germany	X	X	X	X	X	X	Macedonia		X	X		X	
Ghana	X	X	X		X	X	Madagascar	X		X		X	
Greece	X	X	X		X	X	Malawi	X				X	
Grenada					X		Malaysia	X	X	X		X	X
Guam	X		X		X		Maldives			X		X	
Guatemala	X	X	X		X	X	Mali	X	X	X		X	X
Guinea	X		X		X		Malta	X	X			X	X
Guinea Bissau		X	X		X		Mauritania	X		X		X	
Guyana	X				X		Mauritius	X	X	X		X	X
Haiti	X	X	X		X	X	Mexico	X	X	X	X	X	X
Honduras	X	X	X		X	X	Micronesia					X	
Hungary	X	X	X	X	X	X	Moldova	X	X	X		X	X
Iceland	X		X	X	X	X	Mongolia	X	X	X		X	X
India	X	X	X	X	X	X	Montenegro			X		X	
Indonesia	X	X	X		X	X	Morocco	X	X	X		X	X
Iran	X		X	X	X	X	Mozambique	X	X	X		X	X
Iraq	X	X			X	X	Myanmar	X	X			X	X
Ireland	X	X	X	X	X	X	Namibia	X	X	X		X	X
Israel	X	X	X	X	X	X	Nepal	X	X	X		X	X
Italy	X	X	X	X	X	X	Netherlands	X	X	X	X	X	X
Jamaica	X	X	X		X	X	NewZealand	X	X	X		X	X
Japan	X	X	X	X	X	X	Nicaragua	X	X	X		X	X
Jordan	X	X	X		X	X	Niger	X	X	X		X	X
Kazakhstan	X	X	X		X	X	Nigeria	X	X	X	X	X	X
Kenya	X	X	X	X	X	X	Norway	X	X	X		X	X
Kiribati	X				X		Oman	X	X	X		X	X
Kuwait	X	X			X	X	Pakistan	X	X	X	X	X	X
Kyrgyz Rep.	X	X	X		X	X	Panama	X	X		X	X	X
Lao PDR	X	X			X		Papua New Guinea	X	X			X	X
Latvia	X	X	X		X	X	Paraguay	X	X	X		X	X
Lebanon	X	X	X		X	X	Peru	X	X	X	X	X	X
Lesotho	X				X		Philippines	X	X	X		X	X
Poland	X	X	X		X	X	Sweden	X	X	X	X	X	X
Portugal	X	X	X	X	X	X	Switzerland	X	X	X	X	X	X
Qatar	X	X			X	X	Tajikistan	X	X			X	X
Romania	X	X	X		X	X	Tanzania	X	X	X		X	X
Russian Federation	X	X	X	X	X	X	Thailand	X	X	X		X	X
Rwanda	X	X	X		X	X	The Bahamas	X	X	X		X	X
Samoa					X		The Gambia	X		X		X	
Sao Tome & Principe					X		Timor Leste	X		X		X	
Saudi Arabia	X	X	X	X	X	X	Togo	X	X	X		X	
Senegal	X	X	X		X	X	Tonga	X				X	
Serbia	X	X	X		X	X	Trinidad and Tobago	X	X			X	X
Seychelles	X				X		Tunisia	X				X	X
Sierra Leone	X		X		X		Turkey	X	X	X		X	X
Singapore	X	X	X		X	X	Turkmenistan	X				X	
Slovakia	X	X	X		X	X	Uganda	X	X	X		X	X
Slovenia	X	X	X		X	X	Ukraine	X	X	X		X	X
Solomon Islands	X				X		United Arab Emirates	X	X	X		X	X
South Africa	X	X	X		X	X	United Kingdom	X	X	X	X	X	X
South Korea	X	X	X		X	X	United States	X	X	X	X	X	X
Spain	X	X	X	X	X	X	Uruguay	X	X			X	X
Sri Lanka	X	X	X		X	X	Uzbekistan	X		X		X	
St Lucia					X		Vanuatu	X				X	
St Vincent & Grenadines					X		Vietnam	X	X	X		X	X
Sudan	X		X		X		Zambia	X	X			X	X
Suriname	X				X		Zimbabwe	X	X			X	X

The table indicates whether data for a given country on a specific indicator is available and, as indicated in the rightmost column, included in our analysis. 'Stringency' refers to data on containment measures compiled by (Hale et al., 2021). 'Mobility' refers to data on workplace attendance from Google LLC (2021). 'Income' refers to GDP/capita measures from the World Bank World Development Indicators. 'Seroprevalence' refers to the SeroTracker dataset compiled by Arora et al. (2021). 'Deaths' refer to the excess mortality estimates compiled for 191 countries by Wang et al. (2022)

References

Alon, T., Kim, M., Lagakos, D., VanVuren, M., 2020. How should policy responses to the Covid-19 pandemic differ in the developing world? CEPR Covid Econ. (22).

Alvarez, F., Argente, D., Lippi, F., 2021. A simple planning problem for COVID-19 lock-down, testing and tracing. AER: Insights 3 (3), 367–382. http://dx.doi.org/10.1257/aeri.20200201.

Arora, R.K., Joseph, A., Van Wyk, J., Rocco, S., Atmaja, A., May, E., Yan, T., Bobrovitz, N., Chevrier, J., Cheng, M.P., et al., 2021. SeroTracker: a global SARS-CoV-2 seroprevalence dashboard. Lancet Infect. Dis. 21 (4), e75—-e76.

Barnett-Howell, Z., Mobarak, M., Watson, O.J., 2021. The benefits and costs of social distancing in high- and low-income countries. Trans. R. Soc. Trop. Med. Hyg. 115 (7), 807–819. http://dx.doi.org/10.1093/trstmh/traa140, arXiv:2004.04867.

Bick, A., Fuchs-Schuendeln, N., Lagakos, D., 2018. How do hours worked vary with income? Cross-country evidence and implications. Amer. Econ. Rev. 108 (1), http://dx.doi.org/10.1257/aer.20151720.

Boppart, T., Krusell, P., 2020. Labor supply in the past, present, and future: a balanced-growth perspective. J. Polit. Econ. 128 (1), 118–157.

Bryan, G., Chowdhury, S., Mobarak, M., 2014. Underinvestment in a profitable technology: The case of seasonal migration in Bangladesh. Econometrica 82 (5), 1671–1748. http://dx.doi.org/10.3982/ECTA10489.

Charmes, J., 2015. Time use across the world: Findings of a world compilation of time use surveys. UNDP Hum. Dev. Rep. Off..

Chetty, R., Friedman, J.N., Hendren, N., Stepner, M., Opportunity Insights Team, 2021. The economic impacts of COVID-19: Evidence from a new public database built using private sector data. NBER Work. Pap. 26463.

Chudik, A., Pesaran, M.H., Rebucci, A., 2022. Social distancing, vaccination and evolution of COVID-19 transmission rates in europe. IMF Econ. Rev. http://dx.doi.org/10.1057/s41308-022-00181-9.

Desai, S., Pramanik, S., 2020. Virus fear and logistical challenges curtailing economic activity in Delhi-NCR, new survey shows. Livemint URL.

Duarte, M., Restuccia, D., 2010. The role of the structural transformation in aggregate productivity. Q. J. Econ. 125 (1), 129–173. http://dx.doi.org/10.1162/qjec.2010.125.1.129.

Eichenbaum, M., Rebelo, S., Trabandt, M., 2021. The macroeconomics of pandemics. Rev. Financ. Stud. 34 (11), 5149–5187. http://dx.doi.org/10.1093/rfs/hhab040.

Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2022. Inequality in life and death. IMF Econ. Rev. 70 (1), 68–104. http://dx.doi.org/10.1057/s41308-021-00147-3.

Farag, M., NandaKumar, A.K., Wallack, S., Hodgkin, D., Gaumer, G., Erbil, C., 2012. The income elasticity of health care spending in developing and developed countries. Int. J. Health Care Finance Econ. 12 (2), 145–162. http://dx.doi.org/10.1007/s10754-012-9108-z.

Farboodi, M., Jarosch, G., Shimer, R., 2020. Internal and external effects of social distancing in a pandemic. NBER Work. Pap. 27059.

Garibaldi, P., Moen, E., Pissarides, C., 2020. Modelling contacts and transitions in the SIR epidemics model. CEPR Covid Econ. (5).

Ghisolfi, S., Almås, I., Sandefur, J., von Carnap, T., Bold, T., 2020. Predicted COVID-19 fatality rates based on age, sex, comorbidities and health system capacity. BMJ Glob. Health 5 (e003094), http://dx.doi.org/10.1136/bmjgh-2020-003094.

Glover, A., Heathcote, J., Krueger, D., Ríos-Rull, J.-V., 2021. Health versus wealth: On the distributional effects of controlling a pandemic. Work. Pap..

Google LLC, 2021. Google COVID-19 community mobility reports.. URL https://www.google.com/covid19/mobility/.

Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., Tatlow, H., 2021. A global panel database of pandemic policies (oxford COVID-19 government response tracker). Nat. Hum. Behav. 5 (4), 529–538. http://dx.doi.org/10.1038/s41562-021-01079-8.

Hausmann, R., Schetter, U., 2022. Horrible trade-offs in a pandemic: Poverty, fiscal space, policy, and welfare. World Dev. 153, http://dx.doi.org/10.1016/j.worlddev.2022.105819.

Herrendorf, B., Valentinyi, A., Rogerson, R., 2014. Growth and structural transformation. Handb. Econ. Growth.

IGM Forum, 2020. Policy for the COVID-19 crisis. URL http://www.igmchicago.org/surveys/policy-for-the-covid-19-crisis/.

Kim, M., Alon, T., Lagakos, D., VanVuren, M., 2023. Macroeconomic effects of COVID-19 across the world income distribution. IMF Econ. Rev..

King, R.G., Plosser, C.I., Rebelo, S., 1988. Production, growth and business cycles: I. The basic neoclassical model. J. Monetary Econ. 21 (2–3), 195–232.

Krueger, D., Uhlig, H., Xie, T., 2022. Macroeconomic dynamics and reallocation in an epidemic. Econ. Policy 37 (110), 341–398.

Liu, Y., Gayle, A.A., Wilder-Smith, A., Rocklöv, J., 2020. The reproductive number of COVID-19 is higher compared to sars coronavirus. J. Travel Med. 27 (2), 1–4. http://dx.doi.org/10.1093/jtm/taaa021.

Miguel, E., Mobarak, A.M., 2022. The economics of the Covid-19 pandemic in developing countries. Annu. Rev. Econ. 14 (253–285).

Moll, B., 2020. Lockdowns in SIR models. URL.

OECD, 2012. Mortality Risk Valuation in Environment, Health and Transport Policies. OECD Publishing.

Ohanian, L., Raffo, A., Rogerson, R., 2008. Long-term changes in labor supply and taxes: Evidence from OECD countries, 1956–2004. J. Monetary Econ. 55 (8), 1353–1362.

Ravallion, M., 2020. Pandemic policies in poor places. CGD Note URL https://www.cgdev.org/sites/default/files/pandemics-in-poor-places.pdf.

Ray, D., Subramanian, S., 2020. India's lockdown : An interim report. Indian Econ. Rev. http://dx.doi.org/10.1007/s41775-020-00094-2.

Ray, D., Subramanian, S., Vandewalle, L., 2020. India's lockdown. CEPR Policy Insight (102).

Stone, R., 1954. Linear expenditure systems and demand analysis: an application to the pattern of british demand. Econ. J. 64 (255), 511–527.

The Economist, 2022. The pandemic's true death toll. URL https://www.economist.com/graphic-detail/coronavirus-excess-deaths-estimates.

Toxvaerd, F., 2020. Equilibrium social distancing. Camb. Work. Pap. Econ. 2021.

Viscusi, W.K., Masterman, C.J., 2017. Income elasticities and global values of a statistical life. J. Benefit-Cost Anal. 8 (2), 226–250. http://dx.doi.org/10.1017/bca.2017.12.

Wang, H., Paulson, K.R., Pease, S.A., Watson, S., Comfort, H., Zheng, P., Aravkin, A.Y., Bisignano, C., Barber, R.M., Alam, T., Fuller, J.E., May, E.A., Jones, D.P., Frisch, M.E., Abbafati, C., Adolph, C., Allorant, A., Amlag, J.O., Bang-Jensen, B., Bertolacci, G.J., Bloom, S.S., Carter, A., Castro, E., Chakrabarti, S., Chattopadhyay, J., Cogen, R.M., Collins, J.K., Cooperrider, K., Dai, X., Dangel, W.J., Daoud, F., Dapper, C., Deen, A., Duncan, B.B., Erickson, M., Ewald, S.B., Fedosseeva, T., Ferrari, A.J., Frostad, J.J., Fullman, N., Gallagher, J., Gamkrelidze, A., Guo, G., He, J., Helak, M., Henry, N.J., Hulland, E.N., Huntley, B.M., Kereselidze, M., Lazzar-Atwood, A., LeGrand, K.E., Lindstrom, A., Linebarger, E., Lotufo, P.A., Lozano, R., Magistro, B., Malta, D.C., Månsson, J., Mantilla Herrera, A.M., Marinho, F., Mirkuzie, A.H., Misganaw, A.T., Monasta, L., Naik, P., Nomura, S., O'Brien, E.G., O'Halloran, J.K., Olana, L.T., Ostroff, S.M., Penberthy, L., Reiner Jr, R.C., Reinke, G., Ribeiro, A.L.P., Santomauro, D.F., Schmidt, M.I., Shaw, D.H., Sheena, B.S., Sholokhov, A., Skhvitaridze, N., Sorensen, R.J.D., Spurlock, E.E., Syailendrawati, R., Topor-Madry, R., Troeger, C.E., Walcott, R., Walker, A., Wiysonge, C.S., Worku, N.A., Zigler, B., Pigott, D.M., Naghavi, M., Mokdad, A.H., Lim, S.S., Hay, S.I., Gakidou, E., Murray, C.J.L., 2022. Estimating excess mortality due to the COVID-19 pandemic: a systematic analysis of COVID-19-related mortality, 2020–21. Lancet 399 (10334), 1513–1536. http://dx.doi.org/10.1016/s0140-6736(21)02796-3.