

The Macroeconomics of Pandemics around the World: Lives versus Livelihoods Revisited

Ingvild Almås^a

Tessa Bold^b

Tillmann von Carnap^c

Selene Ghisolfi^d

Justin Sandefur^e

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, validate, 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 becomes more rather than less strict. For reasonable income-elasticities of the value of a statistical life, the model can fully rationalize equally strict or stricter policies in poorer countries.

^aProfessor, Institute for International Economic Studies, Stockholm University and Norwegian School of Economics, Bergen, ingvild.almas@iies.su.se

^bAssociate Professor, Institute for International Economic Studies, Stockholm University. Corresponding author: tessa.bold@iies.su.se

^cPhD Student, Institute for International Economic Studies, Stockholm University, tillmann.voncarnap@iies.su.se

^dAssistant Professor, Catholic University, Milan and LEAP, Bocconi University, Milan, selene.ghisolfi@unicatt.it

^eSenior Fellow, Center for Global Development, Washington D.C., jsandefur@cgdev.org

Acknowledgements: We thank Timo Boppert, Tobias Broer, Raymond Guiteras, Per Krusell, David Lagakos, Erzo F. P. Luttmer, Mahreen Mahmud, Mushfiq Mobarak, Rohini Pande, Debraj Ray, Emma Riley, Jakob Svensson, Lore Vandewalle, William Kip Viscusi, Has van Vlokhoven and seminar participants at the IIES Brown Bag, the TARC workshop on COVID-19 and Developing Countries, the FAIR Workshop on COVID-19 in Africa, the SNS International Policy Talk on COVID-19 and Helsinki GSE for helpful comments. Earlier versions of this paper have been published in the 27th issue of the CEPR COVID Economics series and as CGD Working Paper 555.

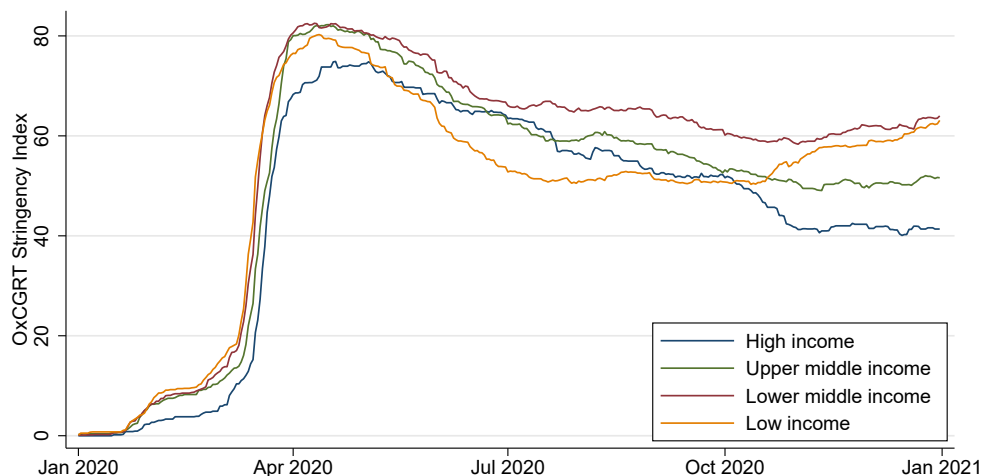
1 Introduction

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 (Figure 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, 2021).

In this paper, we study why containment policy recommendations – by which we refer to a whole suite of non-pharmaceutical interventions, including school closures, travel restrictions, curfews, etc. – might differ for rich and poor countries.¹ To do so, we formalize the role of two distinct factors frequently mentioned in contemporaneous policy debates: i) the lower infection fatality rates anticipated in countries with younger populations, but also frequently weaker health systems and higher incidence of comorbidities, and ii) poverty, which potentially affects the trade-offs made by both economic agents and policymakers between lives and livelihoods.

¹In the analysis that follows we will focus on comparing countries above and below median per capita GDP. Throughout, we refer to these as “rich” and “poor” countries, respectively, and to the difference in their pandemic outcomes as the “rich-poor gap.”

Figure 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.

We develop and estimate a fully specified version of a SIR-macro model, in which agents' and policy-makers' response to the pandemic varies with country-specific measures of fatality risk and income. 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, calibration of the model illustrates that any statement about the desirability of laxer lockdowns in poorer countries is highly sensitive to the value of a statistical life and how it changes with income. For reasonable income-elasticities of the VSL, consistent with existing literature, the model can fully rationalize equally strict or stricter policies in poor countries, despite younger demographics and hence lower infection-fatality rates which unequivocally prescribe laxer policies.

We start from the model of Eichenbaum et al. (2021), henceforth ERT, in which agents expose themselves to infection risk 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 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. To achieve this, we introduce a constant, \bar{c} , in the utility function which makes the marginal rate of substitution between leisure and consumption depend on the level of the latter (King et al., 1988) as commonly found in historical data (Ohanian et al., 2008; Fuchs-Schuendeln et al., 2018; Boppart and Krusell, 2020).

Second, we also allow income levels to affect the trade-off between consumption and catching a potentially deadly infection. Specifically, we allow the value of a statistical life (VSL) to be a function of income and estimate its "income" elasticity using indirect inference. We show that optimal policy depends crucially on this income-elasticity of the VSL as the social planner trades off risks of death against utility of the living.

Third, we also account for the fact that poorer countries face lower infection fatality rates (IFRs) due to their younger demographics (despite higher comorbidities, and weaker health system capacity), which reduce the imperative for lockdowns. We calibrate the IFR using data from Ghisolfi et al. (2020).

To validate the model and estimate underlying behavioral parameters, we use cross-country time-series

data to examine how four key model outcomes – containment policies, economic contraction, COVID-19 infections, and excess mortality – vary across income and IFR. 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.

To assess if our model can rationalize these findings, we simulate the model for 125 countries and a range of parameter values using data on predicted infection fatality rates (IFR) and income (per capita GDP) and generate the same eight statistics that we estimated empirically. We then validate the model and estimate the unknown parameters by indirect inference. Specifically, we consider the model valid only if there are parameters for which the signs of the eight moments match in simulated and observed data. The subset of parameter values for which the signs match produce empirical bounds for the income elasticity of the VSL and the subsistence constraint.

We find that the model successfully matches all eight target moments in the data with an income-elasticity of the VSL which is strictly less than one, and a subsistence constraint that is strictly greater than zero. This 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 estimate 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 that is observed in other contexts.

In terms of the model’s implications, our validation and estimation imply 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., 2021). Similar to our study, Eichenbaum et al. (2022) introduce different VSLs and non-homothetic preferences 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. (2021), 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 SIR (Susceptible - Infected - Removed) 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., 2021). 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., 2021; Krueger et al., 2020; Glover et al., 2020; Alvarez et al., 2021).

In the following section, we lay out our extensions to the basic model, combining epidemiological and economic frameworks. We present comparative statics both analytically and in simulations in Section 3, and calibrate the model for a panel of countries. Section 4 describes our strategy to validate and estimate model parameters via indirect inference. We present empirical estimates of the evolution of excess mortality, infections, lockdowns and economic contraction in Section 5. We estimate and validate the model in Section 6. Section 7 concludes with implications for policy.

2 Model

2.1 The pandemic

We begin with a standard SIRD (Susceptible-Infected-Recovered-Deceased) 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. Firstly, 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 .

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

Secondly, 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).

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

Thirdly, 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 two 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, supported by evidence from the United States, showing that richer households accounted for the bulk of spending reductions during the early stages of the pandemic (Chetty et al., 2020). Second, we make the (otherwise implicit) value of statistical life an explicit function of income, by assuming it displays a constant income-elasticity across countries. We furthermore allow key characteristics of a country that matter for pandemic response—such as the IFR and income—to vary across contexts.

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 discounted lifetime utility as a function of their consumption c_t^j and hours worked n_t^j :

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

The instantaneous utility function depends on consumption c_t and hours worked n_t . We incorporate the subsistence constraint of consumption by including a constant \bar{c} . In general, this constant is assumed to be positive and to represent 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. We will also entertain negative values of \bar{c} below, which allows the elasticity of consumption with respect to a tax to be decreasing in income. Furthermore, 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, which we specify below. With log-preferences for consumption and a convex disutility of labor, this gives

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

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

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

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:

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

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, $c_t^s S_t + c_t^i I_t + c_t^r R_t = C_t$.

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 number of new infections from consumption depends on the exogenous probability of becoming infected when a susceptible person meets an infected one during consumption-related exposure, π_{s1} , times the number of such interactions, given by $(S_t c_t^S)(I_t c_t^I)$. Similarly, the number of new infections at work depends on the exogenous probability of becoming infected when a susceptible meets an infected at work, π_{s2} , times the total hours worked by susceptibles, $S_t n_t^S$ and infected, $I_t n_t^I$. Finally, just as in the standard SIRD, we assume that some share of infections happen irrespective of the intensity of economic interactions, but are simply proportional to the number of infected and susceptible, $\pi_{s3} S_t I_t$. With this, the number of newly infected in period t equals the sum of new infections from the three channels:

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

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

$$\tau_t = \pi_{s1} c_t^S (I_t C_t^I) + \pi_{s2} n_t^S (I_t N_t^I) + \pi_{s3} I_t \tag{2}$$

is the probability of a susceptible becoming infected in period t and 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 (1) by way of a Lagrangian. For the susceptible the maximization involves an additional choice, namely the optimal value of the probability of becoming infected τ_t , and an additional constraint, namely the dependence of this probability the agent's consumption and labor supply (2). 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, $c_t^{s,i,r}$, hours, $n_t^{s,i,r}$, and (for the susceptible) the endogenous probability of infection, τ_t :

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

The additional terms in the first-order conditions for the susceptible, $\lambda_{\tau t} \pi_{s1}(I_t C_t^I)$ in the first-order condition for consumption and $\lambda_{\tau t} \pi_{s2}(I_t n_t^I)$ in the first-order condition 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

containment rate μ_{ct} is chosen by the social planner to maximize the weighted sum of lifetime utilities of the three types of agents in the economy, $W = \sum_t (S_t U_t^s + I_t U_t^i + R_t U_t^r)$, subject to the government’s budget constraint. The optimal containment rate is, in general, non-zero.

3 Model calibration

We now analyze the model’s predictions both in simulations and analytically. In a first step, we illustrate its main mechanisms in our 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.

3.1 Benchmark calibration

For our benchmark case, we follow ERT’s choices for most of the disease and economic parameters (Table 1). Beginning with the disease parameters, ERT cite evidence that about 1/6 of infections take place at work, 1/6 while consuming, and 2/3 during random interactions. We calibrate the π_s using steady state values of hours worked and consumption in the absence of a pandemic, targeting a final epidemic size of 60% and the above infection shares. This yields a basic reproductive number R_0 of 1.5, which is at the lower end of estimates reported for the early stage of the pandemic in Liu et al. (2020).

For the probability of dying from COVID-19 conditional on being infected, we depart from ERT and use country-specific estimates from Ghisolfi et al. (2020), where a country’s IFR is 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%.


We follow ERT in assuming that the average time to either recovery or death is 18 days. For the benchmark case, 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 $\pi_r = \frac{7}{18}(1 - 0.0064)$. We also follow ERT in setting the productivity loss from getting infected at 20% (i.e $\phi = 0.8$).

On the economic side, we need to calibrate the weekly discount factor β , labor productivity A , and the disutility of labor θ . We set the weekly discount factor to $0.96^{\frac{1}{52}}$ and derive the latter two parameters from United States’ pre-pandemic income and hours worked: we convert World Bank annual GDP figures at 2010 constant dollars into weekly values, $Y_{weekly,US} = \frac{55,753}{52}$, set weekly hours for the United States, $H_{weekly,US}$, to 30, and calculate the labor productivity parameter for the United States as $A_{US} = Y_{weekly,US}/H_{weekly,US}$. We find the disutility of labor parameter for the United States by inverting the optimality condition for

Table 1: Mapping model parameters to data

	Definition	Benchmark	All countries	Source
Disease parameters				
π_{sk}	A susceptible's probability of getting infected if exposed during - consumption ($k = 1$) - work ($k = 2$) - other contacts ($k = 3$)	$\pi_{s1} = 7.84 \times 10^{-8}$ $\pi_{s2} = 1.24 \times 10^{-4}$ $\pi_{s3} = 0.3901$	country-specific	ERT
π_d	Prob. of dying within 7 days of infection	$0.0064 \times \frac{7}{18}$	country-specific weekly predicted IFR	Ghisolfi et al (2020); ERT
π_r	Prob. of recovering within 7 days of infection	$(1 - 0.0064) \times \frac{7}{18}$	country-specific	"
ϕ	Productivity while infected	0.8	0.8	ERT
Economic parameters				
β	Annual discount factor	0.96	0.96	ERT
θ	Disutility of labor		from country-specific weekly hours and GDP/capita	World Bank, World Development Indicators
A	Productivity of labor		"	"
\bar{c}	Subsistence constraint	0	$\bar{c} \in [0, 600]$, precise value to be estimated	
\bar{u}	Underpins the VSL	0	$\eta \in [0.7, 1.2]$, precise value to be estimated	Viscusi and Masterman (2017)

Figure 2: Benchmark calibration



The figure shows simulated time series for epidemic and economic outcomes for the benchmark SIR-model (dotted line), the SIR-macro model without policy (solid blue line) and the SIR-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 to model the United States' economy and epidemic.

steady-state hours worked in the absence of a pandemic: $\theta_{US} = \frac{1}{H_{weekly, US}}$.

For our benchmark calibration, we set the value of the newly introduced parameters, \bar{c} and \bar{u} , to zero as in ERT. When $\bar{u} = 0$, the value of a statistical life simply equals the model's discounted lifetime utility of consumption. For this calibration, our utility function, discount factor and target values for consumption and hours worked together imply a VSL for the United States of \$9.3 million and thus a VSL-to-income ratio of 166. This is close to the value used for United States policy making, as reported by Viscusi and Masterman (2017).

Figure 2 shows the results for our benchmark calibration. The dotted black line displays the course of the epidemic without any behavioral adjustment or government intervention – the basic SIRD model. The epidemic cannot end before a sufficient share of the population has acquired immunity – the herd immunity level – calculated as the inverse of the initial reproduction number R_0 and equal to 33%. 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% have ever been infected (targeted by the parameterization), and 0.38% have died (60% infected * 0.64% IFR = 0.38%). Hence, the final epidemic

size as measured by the share of people that ever get 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 model without any containment policy. In response to the pandemic, aggregate consumption (and hours) spontaneously fall by up to 12%, because susceptible agents voluntarily adjust their consumption and labor supply in order to reduce their risk of infection.² The voluntary response reduces the pandemic overshoot by about 10 percentage points or one third and deaths by about 0.8 percentage points or one quarter relative to the standard SIRD.

That the susceptible reduce their hours relative to the steady state, the more so the more severe is the pandemic, can also be seen analytically by inspecting the expressions for optimal hours for the three types of agents $j = s, i, r$

$$n_t^j = -K_t^j + \sqrt{(K_t^j)^2 + \frac{1}{\theta}}, \quad (3)$$

where $K_t^s = \frac{1}{2} \frac{\beta(U_{t+1}^s - U_t^s)}{\theta} (A\pi_{s1}(I_t c_t^I) + \pi_{s2} I_t n_t^I) \geq 0$ and $K_t^{i,r} = 0$. Since

$$\frac{\partial n_t^j}{\partial K_t^j} = -1 + \frac{K_t^j}{\sqrt{(K_t^j)^2 + \frac{1}{\theta}}} = -1 + \frac{\sqrt{(K_j^s)^2}}{\sqrt{(K_t^j)^2 + \frac{1}{\theta}}} < 0,$$

it follows that $n_t^s \leq n_t^j$ for $j = i, r$. Hence, the susceptible reduce their hours relative to the infected and recovered when there is a pandemic, the more so, the larger is K_t^s , which, in turn is increasing in the number of infected, I_t . The voluntary adjustment thus tracks infections and peaks at the same time.

Besides this voluntary adjustment, the social planner can internalize the individual agents' contribution to the pandemic, specifically that of the infected, by taxing consumption and thereby discouraging economic activity beyond the voluntary adjustment. The evolution of the pandemic economy with an optimal 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% during peak infections or 49% on average over the first year. As a result, the economy contracts even further, culminating in a drop of hours and consumption by 35%. 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.

²Mechanically, the infected also reduce their consumption because their labor becomes 20% less productive. This leads only to a small drop in consumption, however: with a peak infection rate of 5%, this amounts to a 1% reduction in consumption.

3.2 Calibrating the full model with heterogeneity in income and the IFR

Having reviewed the basic mechanisms of the model, we now examine how the optimal policy and agents' response to it, change as we move from the benchmark case of the United States to other countries. 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.

In our model, a country's income affects agents' ability and willingness to adjust to disease risk in a number of ways. First, the introduction of \bar{c} 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. Rather, the poorer a country and thus the closer its income to the absolute subsistence level, the less agents will respond to a containment tax.

Second, it is reasonable to suppose that the value of a statistical life varies with income, at least to some extent. Following Viscusi and Masterman (2017), we relate a country's VSL to a base VSL and the country's income relative to the base country via the following function:

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

where η is the common international income elasticity of the VSL. As the base country we use the United States because there are reliable empirical estimates for its VSL. 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 will have a lower (higher) VSL to income ratio than the base case.

To simulate the extended model across countries with a lower observed income and different predicted mortality risk than the base case, we use country-specific values for income, hours worked, and the predicted IFR, while all other parameters are fixed to the same values as in the benchmark case (see Table 1).

Specifically, for each country's weekly income $Y_{weekly,j}$, we take World Bank GDP figures at 2010 constant dollars and convert them into weekly numbers. Following Fuchs-Schuendeln et al. (2018), weekly hours are set to 30 for high and upper middle income countries and to 50 for low and lower middle income countries. With this, we calculate country-specific labor productivity as $A_j = Y_{weekly,j}/H_{weekly,j}$ and the country-specific disutility of labor as $\theta_j = \frac{1}{H_{weekly,j}}$. For country j 's IFR, we use the values predicted by Ghisolfi et al. (2020).

Since there are no widely accepted sets of parameters across countries for the absolute subsistence constraint and the income elasticity of the VSL, we simulate the model on a grid of plausible values for \bar{c} and

η . Later, we will use the model to estimate these parameters.

To allow for varying degrees of non-homotheticity (and a marginal rate of substitution between consumption and leisure that can be both increasing or decreasing in income) we let \bar{c} range from -200\$ to 600\$ measured in yearly consumption.³ 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. In order for the model to produce the desired $VSL_j(\eta)$ as its discounted lifetime steady-state utility, we then adjust \bar{u} accordingly:⁴

$$\bar{u} = VSL \frac{1-\beta}{c-\bar{c}} - \ln(c-\bar{c}) - \frac{\theta}{2}n^2.$$

Note that if \bar{u} was set to zero for each country, the VSL implied by the model would simply equal the sum of the lifetime utility for the living. For poorer countries, this ‘model-implied’ VSL implies an income-elasticity of the VSL higher than the maximum we consider on our grid and a VSL-income ratio much lower than the one for the US. For example, in Uganda, the VSL-income ratio would be 52, one third of the US’ VSL-income ratio.

3.3 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 are monotone functions of the underlying preferences, η and \bar{c} , and vary substantially and systematically with the dimensions of observed heterogeneity, income, and IFR. To do so, we simulate the model’s time series for illustrative scenarios of countries differing by income and IFR. We set pre-pandemic incomes equal to the World Bank thresholds for low-income, lower-middle income and upper-middle income countries, and we calibrate the productivity parameter and disutility of labor parameter for these example cases. The IFR takes on values equal to the 25th, 50th, and 75th percentile of the global predicted distribution in Ghisolfi et al. (2020). For each scenario, the model is solved for the parameter ranges discussed above, namely $\eta = \{0.7, 1.3\}$ and $\bar{c} = \{-200, 600\}$.

The comparative statics are shown in Figure 3. Each row of the figure graphs one of three measures

³Another consideration is that per-capita income of the poorest country must exceed the absolute subsistence constraint in order for per-period utility to be defined.

⁴To see how this is done, write down the VSL as the lifetime utility in consumption units, i.e.

$$\begin{aligned} VSL &= \frac{U}{u'(c)} \\ &= \frac{1}{1-\beta} u(c, n) \times \left(\frac{1}{c-\bar{c}}\right)^{-1} \\ &= \frac{c-\bar{c}}{1-\beta} (\ln(c-\bar{c}) + \frac{\theta}{2}n^2 + \bar{u}) \end{aligned}$$

and solve for \bar{u} .

of the economic response (peak drop in hours worked with no containment tax, peak optimal containment tax, and peak drop in hours with the optimal containment tax) and two measures of the pandemic's evolution (cumulative infections and cumulative deaths from COVID-19) on the vertical axis as function of the parameters on the horizontal axis.⁵ The first rows of panel (a) and panel (b) vary the income-elasticity of the VSL, η , (with $\bar{c} = 0$) and the second row in panel (a) and panel (b) vary the subsistence level, \bar{c} (with $\eta = 1$). Each subgraph in panel (a) contains three curves representing the three income scenarios (GDP/capita=1045\$, yellow line; GDP/capita=4096\$, green line; GDP/capita=12696\$, blue line) with the IFR fixed at the median across countries, and each subgraph in panel (b) contains three curves representing the three IFR scenarios ($IFR = 0.37\%$, yellow line; $IFR = 0.53\%$, green line; $IFR = 0.74\%$, blue line) with income fixed at the median across countries. The lines' slopes thus represent the partial derivative of the given outcome with respect to the parameter on the horizontal axis (at different income and IFR levels). The height difference of the lines, on the other hand, provides a measure of the sign and (approximate) magnitude of the model outcome's partial derivative with respect to income and the IFR (at different levels of η and \bar{c}).

Figure 3 illustrates several points. First, both individuals' labor supply response and the optimal policy are strictly and strongly decreasing in η (see the first rows of panel (a) and panel (b), respectively). For example, with an $\eta = 0.7$ the representative lower-middle income country will set a maximum containment tax of 180%, while with $\eta = 1.3$ the same country will set a maximum containment tax of 55%. Together, these two variables determine the total drop in economic activity, which is therefore also strictly decreasing in η . Consequently, infections and deaths are strictly increasing in η since the larger the economic contraction, the lower are infections and deaths.

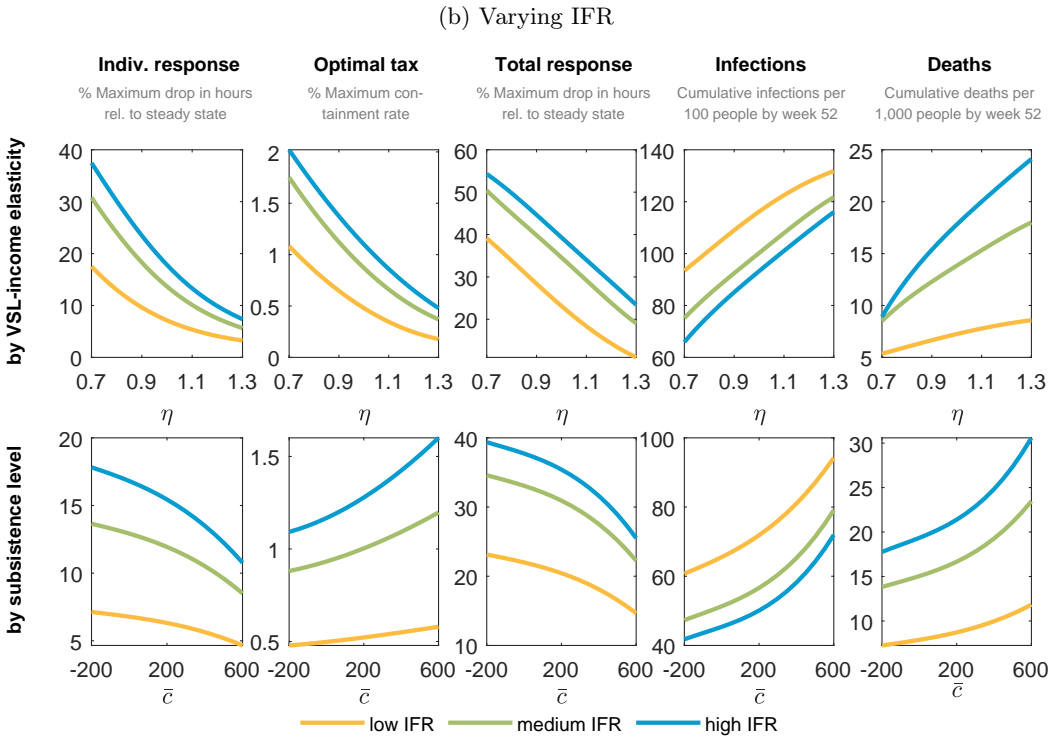
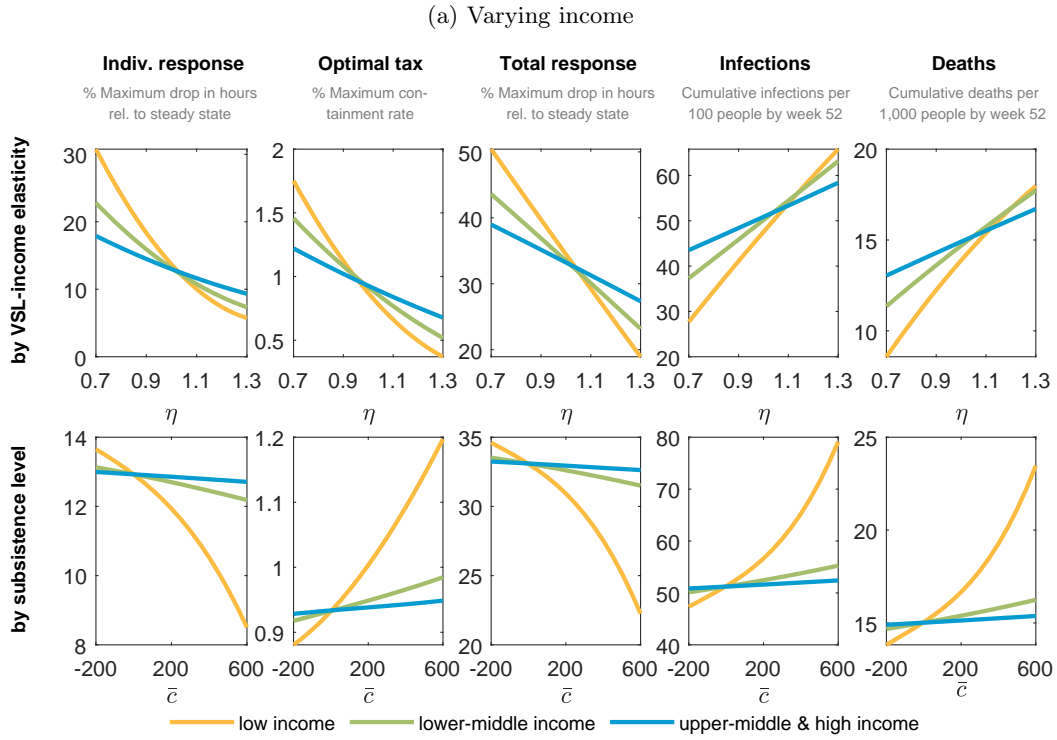
The partial derivatives of the individual response, the optimal containment tax and the total economic contraction with respect to η are negative (and the partial derivative of infection and deaths thus positive) at all income (and IFR) levels, because the incomes we consider are lower than the US benchmark, so the VSL to income ratio is decreasing with η for the representative cases. Intuitively, the higher is η , the lower is the utility cost of dying and so is hence the economic and policy response to the pandemic.

Second, a higher subsistence level makes agents less willing to reduce their consumption and labor supply, both in response to infection risk and in response to a given tax rate. As a result both the individual and the total economic contraction are decreasing in the subsistence constraint (see second row in panel (a) and panel (b)). The optimal policy, however, is increasing in the subsistence level.

Why do the optimal policy and the individual and total economic contraction move in opposite directions?

⁵We compare the maxima and cumulative values of such outcomes through time as these are sufficient statistics for the "strength" of the model's response to the pandemic.

Figure 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.

The standard effect of a positive subsistence constraint is to make all agents respond less to the pandemic. That is, the infected increase their labor supply relative to the recovered in order to compensate for the productivity drop that arises as a consequence of infection. For the susceptible, 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} . The policymaker is now faced with a world where agents find a reduction in hours and consumption more costly, 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 that is decreasing in the subsistence level. With the total drop in economic activity decreasing in the subsistence level, the model predicts worse epidemic outcomes, i.e. higher infections and deaths, the higher the subsistence constraint.

Third, the optimal containment policy and the total contraction in economic activity are both strictly increasing in mortality risk at all levels of η and \bar{c} (see panel (b)).⁶ As a result, infections, which are more costly the higher the IFR, are decreasing in mortality risk. In the end, this is not enough for high IFR countries to reduce deaths below those of low IFR countries though, and deaths remain increasing in the IFR even in the SIR-Macro model with optimal policy.

Hence, the model does not appear to overturn the intuition that the lower IFR of poor countries should drive them towards laxer policies. The picture is very different, however, for the other dimension of observed heterogeneity, pre-pandemic income. The partial derivative of the model's outcomes with respect to income can take on both large positive values, large negative values and everything in between (see the first row of panel (a) and panel (b), respectively).

Specifically, and the fourth point worth highlighting in Figure 3, the effect of income on the individual response, optimal containment and total economic contraction is negative for values of $\eta < 1$, zero at $\eta = 1$, and positive for $\eta > 1$. The reason is that agents and policy-makers in countries that are poorer than the benchmark case of the United States will have a higher VSL-income ratio than the US when $\eta < 1$ and will thus do more to avoid fatality risk. The opposite holds when $\eta > 1$. Consequently, infections and deaths are decreasing in income when $\eta < 1$ but increasing when $\eta > 1$. At $\eta = 1$, (and $\bar{c} = 0$) all outcomes are independent of income since the VSL-income ratio is constant across income levels.

Fifth, poor countries exhibit a more muted individual response and total economic contraction for values

⁶This can be seen analytically from the optimality condition of hours worked for the susceptible in (3): the utility difference between being susceptible and becoming infected $U_t^s - U_t^i$ is increasing in the probability of dying from COVID-19, and hours are thus decreasing in the IFR.

of $\bar{c} > 0$ than rich countries. The reason is that for a positive subsistence constraint, poorer countries shield less and respond less strongly to a tax (i.e. the partial derivative of these outcomes wrt income is positive). The converse holds for the optimal containment tax: since individual shielding and the response to a given tax is weaker in poorer countries when \bar{c} is positive, the optimal tax is higher. Infections and deaths follow the total contraction in economic activity, which are thus higher in poorer countries when $\bar{c} > 0$ (i.e. the partial derivative of infections and deaths wrt income is negative).

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.

The comparative statics show that the intuition that low-income countries should set laxer policies is strongly dependent on the underlying preferences. In particular, even when explicitly modeling the notion that poor countries must trade-off deaths from COVID-19 against deaths from poverty when choosing lockdowns, we do not find that poorer countries should set laxer policies. On the contrary, since a subsistence constraint leads to more infections as both infected and susceptible shield less and respond less to a given tax, the policymaker sets a stricter policy. Similarly, the model predicts that poorer countries set stricter policies than richer ones when the income-elasticity of the VSL is less than unity.

4 Model estimation

The comparative statics exercise has shown the basic mechanisms of a macro-pandemic model with heterogeneity in income, the VSL (parameterized by a constant income-elasticity), the sensitivity of consumption to a tax at different income levels (parameterized by an absolute subsistence constraint), and the IFR. Specifically, the model describes the evolution of the endogenous outcomes \mathbf{Y}_t (individual response, optimal containment tax, contraction of the economy, infections and deaths) as a function of their lags, of exogenous variables (pre-pandemic income level and the IFR), and of parameters (η, \bar{c}) ,

$$\mathbf{Y}_t = f(\mathbf{Y}_{t-1}, \text{Income}, \text{IFR}; \eta, \bar{c}).^7$$

The overarching question is whether the similarity of policy across rich and poor countries implies 'unreasonable' trade-offs between lives and livelihoods in poor countries. To answer this question, we must estimate the unknown parameters that drive these trade-offs, namely the income-elasticity of the VSL and the subsistence constraint, and the macro-pandemic model provides the means to translate observed policy

choices into these unobserved preferences. Since the model does not deliver an analytic expression for the evolution of the endogenous outcomes, we use simulation-based indirect inference (see Smith (2008) and Gouriéroux and Monfort (1997)) to estimate the model’s unknown parameters and assess its goodness of fit with country-aggregate data from 125 countries over the course of the pandemic.

Briefly, we proceed by first simulating for a given parameter combination (e.g., η and \bar{c}) the model’s endogenous outcomes for the joint distribution of the observed exogenous variables (pre-pandemic income and IFR in 125 countries). In the second step, we choose ‘target’ moments that summarize relevant aspects of the joint distribution of the simulated sequences and that we aim to match in the observed data (namely, the gap in pandemic outcomes between rich and poor countries, and between low- and high-IFR countries). We calculate these moments for each parameter combination, and give a heuristic argument how they identify the unknown parameters and validate of the model. We make inference about the unknown parameters by choosing those values of the parameters that make the simulated target moments match those same moments estimated in data. 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), we focus on matching the signs of the target moments, not their magnitudes, in simulated and observed data. We consider the model empirically valid if the sign pattern of the 8 target moments in the data can be generated by the model.

4.1 Simulated moments


The first task is to create the simulated data, from which the model’s summary moments will be calculated. To this end, we numerically solve the model for each of 125 countries using the country’s observed pre-pandemic GDP/capita, hours worked, and predicted IFR as the exogenous inputs. For each country, the model is solved on a 7×5 grid for the unknown parameters η and \bar{c} with η ranging from 0.7 to 1.3 in steps of 0.1 and \bar{c} ranging from -200 to 600 in steps of 200. For each parameter combination on this grid, we create simulated sequences for optimal policy, the drop in economic activity before policy intervention, the total drop in economic activity, infections, and deaths for the first 100 weeks of the pandemic, for a total of 35 sets of simulated data for 125 countries on the whole parameter space.

Next, we need to decide on target moments that will be used to identify the unknown parameters and assess the model’s validity by comparing the simulated moments across the parameter space with the same moments in the observed data. To estimate the two unknown parameters, we need at least two target moments (in order to formulate two moment conditions), but to validate the model we need more moments so the model is over-identified.

Informed by the analysis in Section 3.3, we choose eight target moments, which are the expectations of peak policy (i.e. the strictest policy observed over the first year of the pandemic), peak contraction in economic activity, peak infections, and peak cumulative deaths, conditional on the country being above or below median income and the expectations of those same four variables conditional on the country being above or below the median IFR in the global distribution. We focus on four of the five endogenous outcomes which were used in the comparative statics, omitting the spontaneous contraction in economic activity by individuals as it is not directly observable. For each variable, we let the estimation target be on peak and cumulative magnitudes rather than the whole time series to adjust for the different timing of the pandemic across countries. We focus on differences in moments across the dimensions of observed heterogeneity rather than levels, because differences filter out all unmodelled heterogeneity that is uncorrelated with the IFR and income.

In Figure 4, 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 Figure 4, 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).

Figure 4: 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)

Relative to the comparative statics in the previous section, the qualitative model predictions remain the same when plotting the model moments on the full two-dimensional parameter space and for the full joint distribution of income and IFR in the sample. That is, 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. Similarly, it remains the case, that rich-poor gaps in policy, economic contraction, infections and deaths can take on both large positive, large negative values and anything in between depending on the parameters. And just as before, the rich-poor policy gap is increasing in η and decreasing in \bar{c} (indicated by all the contour lines being upward sloping). Similarly, the rich-poor gap in the economic contraction, infections and deaths is increasing in both η and \bar{c} (indicated by all the contour lines being downward sloping).

Unlike the comparative statics, the contour lines make explicit the trade-off between η and \bar{c} to create a given outcome. For example, the contour lines for optimal policy show that an income elasticity of the VSL of $\eta = 0.72$ and a subsistence level of $\bar{c} = 0$ will produce the same rich-poor policy gap as an $\eta = 0.85$ and a $\bar{c} = 600$. In the first case, policymakers in poor countries have a much higher VSL to income ratio than policymakers in rich countries, but have no need to counteract limited voluntary shielding by the susceptible. In the second case, the difference in the VSL-income ratio between poor and rich countries is still positive but smaller, but the poor country policymaker reacts to a lower individual response and sensitivity to a given tax by setting an even stricter policy. These two effects cancel each other out to produce the same policy gap.

Another important feature of the contour map relative to the two-dimensional analysis in Section 3.3 is that 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% higher chance of having an IFR below the median. As 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.

4.2 How do the moments identify the parameters and validate the model?

With the aid of the contour map, we now give a heuristic argument for how heterogeneity in the target moments across the observed dimensions can be used to make inference about heterogeneity along the unobserved dimensions and to validate the model.

It is clear from the contour map that the gap in policy across countries with different income levels contains a lot of information about the unknown parameters, especially η : when the rich-poor gap takes on positive values, we can infer that η is above, roughly, 0.9 and when the rich-poor gap takes on negative values, we can infer that η is below 0.9. When rich and poor countries set the same policy (or, for that matter, for any given policy gap), the gap in the total economic contraction is directly informative about the subsistence constraint.

Together variation in these two moments is thus sufficient to identify the unknown parameters. As a concrete example, a zero policy gap is consistent with either $\eta = 0.905$ and $\bar{c} = 0$, with $\eta = 0.97$ and $\bar{c} = 600$ or with any combination of η and \bar{c} on the curve joining those two parameter combinations. This indeterminacy disappears, however, when we look at the policy gap and the economic contraction gap simultaneously. If poor and rich countries set the same policy and poor countries experience a larger economic contraction, then we know that the subsistence constraint is low (and thus η must be low). On the other hand, if poor countries experience a smaller drop in economic activity despite setting the same policy, then we know that the subsistence constraint must be high (and thus so is η). Thus, considering the poor-rich gaps in policy and economic contraction in combination can lead to identification.

While two moment conditions are necessary to estimate the unknown preference parameter, they provide too little information about the likely validity of the model. For example, if the two simulated target moments spanned all of \mathbb{R} across the parameter space, then there would always exist a parameter combination for which the model can generate the target moments observed in the data. Thus to assess the model's validity, further over-identifying restrictions are needed, which open up the possibility that the model can match the data for some values of the target moments observed in data, but not for others. If there are no parameters for which the simulated moments match the estimated moments, this provides evidence against the model as the data generating process.

The policy and economic contraction gaps across high- and low-IFR countries are ideally suited to validate of the model. First, these moments relate to the mechanism that is the hallmark of macro-pandemic models, namely that the economic choices of individual agents and the policy-maker during a pandemic are influenced by disease parameters. Second, the contour map shows that regardless of the unobserved preferences, both gaps are always positive (even at low levels of the income elasticity and high subsistence constraints), i.e.

the model predicts that high IFR countries always set a stricter policy and experience a stronger economic contraction than low IFR countries. Thus, if we were to find the opposite in data, i.e. policies and/or economic contraction are stronger in low IFR countries, this would present strong evidence against a macro-pandemic model.

Finally, the difference in infections and deaths across rich and poor countries and across high- and low-IFR countries provide four more over-identifying restrictions. In some sense, one might argue that these moments do not provide much additional information beyond the policy and economic contraction moments. That is, regardless of the data generating process, it seems reasonable that, all else equal, the more economic activity is reduced, the lower are infections. Thus, once one has matched the economic contraction moments in data and model, one should also be able to match the infection moments (regardless of the data generating process). Similarly, it must be true, all else equal, that deaths are the product of infections times the mortality rate. Thus, once one has matched infections, one should also be able to match deaths (again, regardless of the data generating process).

This discussion suggests that these four target moments provide information about the ‘all else equal’ assumption of our model, and are informative about whether our model adequately describes the data generated in the time of the pandemic. To illustrate, if we found that deaths are higher in low-IFR countries than in high-IFR countries or that both the total economic contraction and infections are larger in rich countries, this would indicate relevant unobserved heterogeneity that the model does not consider.⁸

4.3 Distance metric

Under the assumption that the empirical proxies are positively correlated with the model outcomes (both unconditionally and conditional on income and IFR), it can be shown that the model predicts the same signs for the target moments for model outcomes and proxy variables, even though the magnitudes may differ. That is, the assumption guarantees that if the model predicts stricter policy among high-income countries for a given parameter combination, it also predicts a stricter lockdown for this parameter combination. Thus, we define a criterion to evaluate the discrepancy between the simulated and estimated moments at each point in the parameter space that depends only on the differences in their signs.

Specifically, we choose parameters to minimize the sum of the squared distances between the signs of the simulated and the estimated target moments

$$(\eta, \bar{c}) = \arg \min \sum_X \sum_Y (\mathbb{I}(\Delta_X \tilde{Y}(\eta, \bar{c}) \geq 0) - \mathbb{I}(\Delta_X \hat{Y} \geq 0))^2 \quad (4)$$

⁸An example of a potential heterogeneity, is the COVID-19 transmission rate from work or consumption activities that could vary systematically across rich and poor countries, but the model assumes that it is not correlated with income.

where \mathbb{I} is an indicator function equal to 1 if the target moment (in either simulated or observed data) is greater than or equal zero, and equal to zero otherwise. In words, the criterion function at η , \bar{c} equals the number of gaps of policy, economic contraction, infections, and mortality between high- versus low-IFR countries and between rich versus poor countries that do not have the same signs in simulated and observed data. It attains its minimum of zero at parameter values, for which the signs of all eight target moments are the same in simulated and observed data.

The value of the minimized criterion function can also be used to assess the model’s validity. The fewer signs of the estimated target moments the model is able to match, the larger this value will be and, intuitively, the more skeptical we should be that the model generated the data. In what follows, we want to err on the side of caution and will only use the model to estimate the unknown parameters, if it can match all the signs of the estimated moments and thus the minimized criterion function is zero.

Figure 5 shows what information we can gain about the unknown parameters and the model’s validity from minimizing the distance criterion. In Panel A, we superimpose on one graph the zero-level curves for each of the eight moments in Figure 4. Since the zero-level curves for the IFR gaps lie outside the parameter space, the figure contains only four zero-level curves for the rich-poor gaps, each of which partitions the parameter space into two regions, which are classified according to whether they produce positive or negative gaps 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 A to G, one for each sign pattern that the model is able to generate. From the contour maps, the IFR gaps are $\Delta P > 0$, $\Delta ED > 0$, $\Delta I < 0$ and $\Delta D > 0$ in each of these regions, while the signs of the rich-poor gaps vary across the regions and can be read off from the sign classifications of each region. For example, region D is classified as $\Delta P < 0$, $\Delta ED > 0$, $\Delta I < 0$ and $\Delta D > 0$ for the rich-poor gaps. Thus, all parameter combinations in region D generate the sign pattern $\{+, +, -, +, -, +, -, +\}$ for the simulated target moments, where the first four signs refer to the IFR gaps and the last four refer to the rich-poor gaps. Conversely, if we observed this sign pattern in the data, the set of minimizers of the criterion function would equal region D.

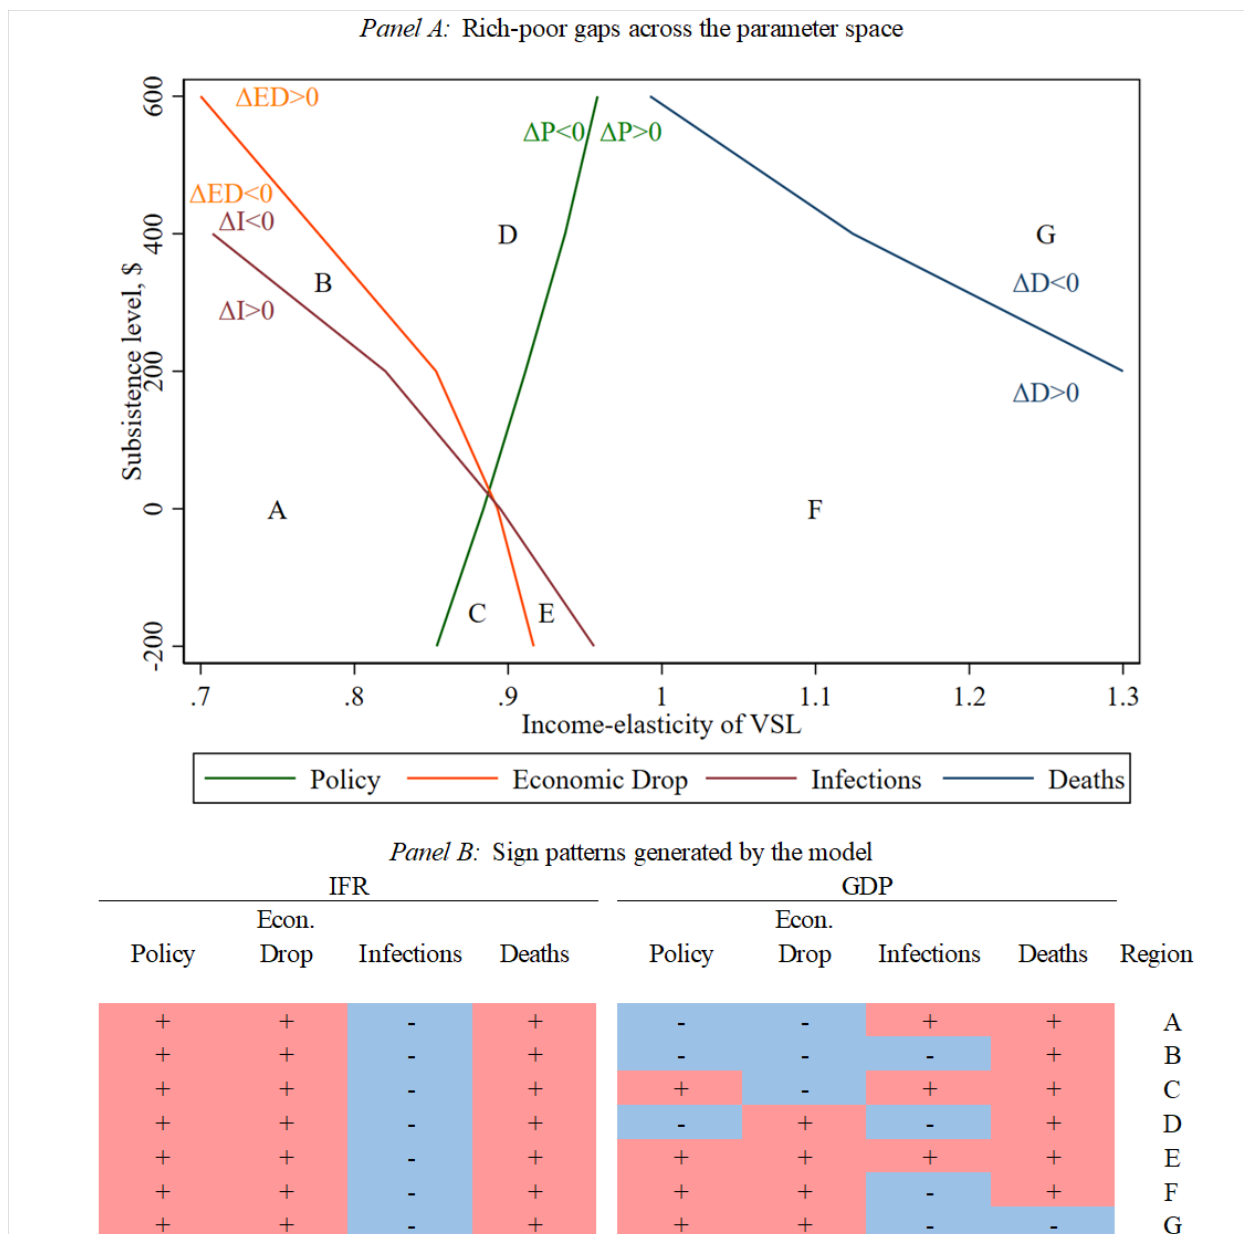
In Panel B, we read off the seven sign patterns that correspond to the individual regions in Panel A of the Figure. As stated, we use the model for estimation only if we observe one of these seven sign patterns in the data.

5 Empirical evidence

We now turn to the empirical counterparts of the simulated moments presented in the previous sections, and show how actual containment policies, economic contractions, infections, and deaths varied between rich and

poor as well as high- and low-IFR countries. 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. We then use the observed differences to (i) validate the model, (ii) estimate the unobserved preference parameters (η and \bar{c}) using the indirect inference method. We conclude by discussing

Figure 5: Set-identification of the parameters



Panel A combines the zero-level curves for each of the eight moments in Figure 4. 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 an either positive or 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 A to G. Each region identifies one sign pattern the model can generate. In Panel B, we read off the seven sign patterns that correspond to the regions in Panel A.

the plausibility of the estimated parameters.

5.1 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. It varies 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, 2021), which tracks how 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 2020.⁹ As an alternative indicator, we also compute differences in per capita consumption from national accounts data for 2020, available at the annual level 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.¹⁰

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.,

⁹It 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., 2020), 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.

¹⁰We confirm that the timing of the seroprevalence studies does not differ between high and low income countries.

Hale et al. (2021)), there is strong evidence they severely under count 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).¹¹ The result is an estimate of cumulative excess mortality at the end of the first year of the pandemic for 174 countries.

5.2 Empirical results

We now estimate the gap in outcomes between high and low-IFR countries and rich and poor countries by regressing each outcome on an indicator variable for whether a country falls above or below the global median for either the predicted IFR or for income per capita. 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.

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% in high-IFR countries, while these numbers stood at the 2.00 and 83% 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% in high-IFR countries and by 49% 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% in high-IFR countries and 17% 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

¹¹As 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

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.

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% in both sets of countries. Despite similar containment, the economy contracted 16% 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% compared to 9%.

6 Indirect inference

The data show that 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. Consistent with the model predictions in Section 2, countries with lower IFRs set laxer policies, experienced less of an economic contraction, and had higher infections, and lower death rates.

6.1 Preference Estimation

We now minimize the distance criterion in (4) to assess if policymakers behave consistently with our model, and to estimate the preference parameters. Based on the regression results in Table 2, the target moments estimated in the data deliver the sign pattern $\{+, +, -, +, -, +, -, +\}$.


Comparing this pattern to the seven patterns in Panel B of Figure 5, we see that it indeed equals one of the seven sign patterns the model can generate. Specifically, the model can generate the observed sign pattern for parameter combinations in region D. Thus, the criterion function is minimized for any parameter combination in D, and its minimum value is 0. Hence, we consider the model validated and estimate lower and upper bounds for η of 0.75 and 0.9 and for \bar{c} of \$50 and \$600.

Do these estimates imply 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 may put a relatively higher weight on lives over 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.

Figure 6: Alignment between rich-poor differences from model and data across outcomes



The figure shows ratios between rich and poor countries’ modelled (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 rich and poor countries are approximately equal. A difference between blue and red bars implies that the ratios are different between modelled 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 (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).

6.2 Comparing calibrated model and data

Having used the congruence of signs to validate the model and 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 countries estimated 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 ratio between poor and rich countries’ outcomes ($\bar{Y}_{poor}/\bar{Y}_{rich}$) and compare each ratio to its data counterpart.

Figure 6 shows that the ratios are on the same side of unity for both model and data in all four cases. Moreover, the ratios line up closely also when looking at the magnitudes, with the exception of the infection data, where the actual ratio of infection rates in poor countries compared to rich countries is much higher than our model predicts. If anything, this discrepancy makes the policy response of low-income countries appear even more reasonable.

Overall, the broad congruence between model predictions and data for rich and poor countries across these four outcomes makes us confident that even a relatively parsimonious model such as ours can be used to analyze pandemic policy-making.

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 both validate the model and estimate underlying 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 with varying levels of the value of a statistical life and 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 containment measures to a rich country, despite – and in fact, because – agents in the poorer context face subsistence constraints that make it harder for them to comply with containment policies.

References

- Alon, T., Kim, M., Lagakos, D., VanVuren, M., 2020. How Should Policy Responses to the COVID-19 Pandemic Differ in the Developing World? NBER Working Paper 27273. doi:10.3386/w27273.
- Alvarez, F., Argente, D., Lippi, F., 2021. A simple planning problem for COVID-19 lock-down, testing, and tracing. *American Economic Review: Insights* 3, 367–82. doi: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. *The Lancet Infectious Diseases* 21, e75–e76.
- Barnett-Howell, Z., Watson, O.J., Mobarak, A.M., 2021. The benefits and costs of social distancing in high- and low-income countries. *Transactions of The Royal Society of Tropical Medicine and Hygiene* doi:10.1093/trstmh/traa140.
- Boppart, T., Krusell, P., 2020. Labor supply in the past, present, and future: a balanced-growth perspective. *Journal of Political Economy* 128, 118–157.
- Chetty, R., Friedman, J., Hendren, N., Stepner, M., Team, T.O.I., 2020. The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data.
- Chudik, A., Pesaran, M.H., Rebucci, A., 2021. COVID-19 Time-varying Reproduction Numbers Worldwide: An Empirical Analysis of Mandatory and Voluntary Social Distancing. *Globalization Institute Working Paper* 407. doi:https://doi.org/10.24149/gwp407.
- Desai, S., Pramanik, S., 2020. Virus fear and logistical challenges curtailing economic activity in Delhi-NCR, new survey shows. *livemint* URL: bit.ly/livemint-dcvts.
- Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2021. The Macroeconomics of Epidemics. *The Review of Financial Studies* 34, 5149–5187. doi:10.1093/rfs/hhab040.
- Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2022. Inequality in Life and Death. *IMF Economic Review* 70. doi:https://doi.org/10.1057/s41308-021-00147-3.
- Farang, 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. *International Journal of Health Care Finance and Economics* 12, 145–162. doi:10.1007/s10754-012-9108-z.

- Farboodi, M., Jarosch, G., Shimer, R., 2021. Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory* 196. doi:10.1016/j.jet.2021.105293.
- Fuchs-Schuendeln, N., Lagakos, D., Bick, A., 2018. How do hours worked vary with income? Cross-country evidence and implications. *American Economic Review* 108. doi:10.1257/aer.20151720.
- Garibaldi, P., Moen, E., Pissarides, C., 2020. Modelling contacts and transitions in the SIR epidemics model. *CEPR Covid Economics* 5.
- Ghisolfi, S., Almås, I., Sandefur, J., von Carnap, T., Bold, T., 2020. Calculating regional differences in COVID-19 fatality rates. *BMJ Global Health* 5. doi:10.1136/bmjgh-2020-003094.
- Glover, A., Heathcote, J., Krueger, D., Ríos-Rull, J.V., 2020. Health versus Wealth: On the Distributional Effects of Controlling a Pandemic. *CEPR Covid Economics* 6.
- Google, 2021. Google COVID-19 Community Mobility Reports. URL: <https://www.google.com/covid19/mobility/>.
- Gouriéroux, C., Monfort, A., 1997. *Simulation-based Econometric Methods*. doi:10.1093/0198774753.001.0001.
- 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). *Nature Human Behaviour* 5, 529–538. doi:10.1038/s41562-021-01079-8.
- Hausmann, R., Schetter, U., 2022. Horrible trade-offs in a pandemic: Poverty, fiscal space, policy, and welfare. *World Development* 153. doi:https://doi.org/10.1016/j.worlddev.2022.105819.
- 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., 2021. Macroeconomic Effects of COVID-19 Across the World Income Distribution. Working Paper .
- King, R.G., Plosser, C.I., Rebelo, S.T., 1988. Production, growth and business cycles: I. The basic neoclassical model. *Journal of Monetary Economics* 21, 195–232.
- Krueger, D., Uhlig, H., Xie, T., 2020. Macroeconomic dynamics and reallocation in an epidemic. *CEPR Covid Economics* 5.

Liu, Y., Gayle, A.A., Wilder-Smith, A., Rocklöv, J., 2020. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine* 27, 1–4. doi:10.1093/jtm/taaa021.

Miguel, E., Mobarak, A.M., 2021. The economics of the Covid-19 pandemic in developing countries. NBER Working Paper 29339.

Moll, B., 2020. Lockdowns in SIR Models. URL: https://benjaminmoll.com/wp-content/uploads/2020/05/SIR_notes.pdf

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. *Journal of Monetary Economics* 55, 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 Economic Review* 55. doi:10.1007/s41775-020-00094-2.

Ray, D., Subramanian, S., Vandewalle, L., 2020. India's Lockdown. CEPR Policy Insight .

Smith, T., 2008. Indirect Inference, in: *The New Palgrave Dictionary of Economics*. 2 ed.

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. *CEPR Covid Economics* 15.

Viscusi, W., Masterman, C.J., 2017. Income Elasticities and Global Values of a Statistical Life. *Journal of Benefit-Cost Analysis* 8, 226–250. doi: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., Alorant, 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., Fedosseva, 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., Spurlack, 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. *The Lancet* 399, 1513–1536. doi:10.1016/s0140-6736(21)02796-3.

A Appendix

Table A1: Countries included in analyses

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 (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)

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	
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

Continued on next page

Table A1, continued

Country	Stringency	Mobility	Income	Seroprevalence	Deaths	Model	Country	Stringency	Mobility	Income	Seroprevalence	Deaths	Model
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

Continued on next page

Table A1, continued

Country	Stringency	Mobility	Income	Seroprevalence	Deaths	Model	Country	Stringency	Mobility	Income	Seroprevalence	Deaths	Model
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