

THE INTERGENERATIONAL MORTALITY TRADE-OFF OF COVID-19 LOCKDOWN POLICIES*

BY LIN MA, GIL SHAPIRA, DAMIEN DE WALQUE, QUY-TOAN DO, JED FRIEDMAN,
AND ANDREI A. LEVCHENKO

*Singapore Management University, Singapore; World Bank; University of Michigan, NBER,
and CEPR*

In lower-income countries, the economic contractions that accompany lockdowns to contain COVID-19 transmission can increase child mortality, counteracting the mortality reductions achieved by the lockdown. To formalize and quantify this effect, we build a macrosusceptible-infected-recovered model that features heterogeneous agents and a country-group-specific relationship between economic downturns and child mortality and calibrate it to data for 85 countries across all income levels. We find that in some low-income countries, a lockdown can produce net increases in mortality. The optimal lockdown that maximizes the present value of aggregate social welfare is shorter and milder in poorer countries than in rich ones.

1. INTRODUCTION

Governments across the world introduced unprecedented lockdown policies in an attempt to contain the spread of COVID-19. Unsurprisingly, a debate soon erupted on what type of lockdowns were warranted and whether the benefits of such policies justify the accompanying dramatic economic contractions. Embracing utilitarianism, economists, among others, focused on the trade-off between the lives saved by a lockdown and its economic costs (Hall et al., 2020; Kim and Loayza, 2021). On the other hand, both proponents of deontological ethics and critics of the statistical value of life recused a policy analytic approach that involves the monetary valuation of life (Viscusi and Aldy, 2003; Slovic and Peters, 2006; Singer and Plant, 2020). This article casts a new light on this debate by uncovering and quantifying an intergenerational mortality trade-off inherent to pandemic mitigation: the disease and the lockdown policies affect the mortality of younger and older individuals differentially.

In the early days of the pandemic, evidence emerged that the COVID-19 mortality risk increases substantially with age (Verity et al., 2020). On the other hand, empirical evidence has shown that infant and child mortality in low- and middle-income countries is countercyclical (Pritchett and Summers, 1996; Bhalotra, 2010; Baird et al., 2011; Cruces et al., 2012; Friedman and Schady, 2013). This implies that lockdown policies in developing countries can lead to an increase in infant and child mortality due to the consequent economic contraction. Thus, pandemic mitigation policies in low-income settings not only forgo economic well-being to save lives but also embed a trade-off between one life and another.

*Manuscript received June 2021; revised January 2022.

We are grateful to the editor (Jesús Fernández-Villaverde), three anonymous referees, Mohamed Abdel Jelil, Daron Acemoglu, Jishnu Das, Shanta Devarajan, Xavier Devictor, Rema Hanna, Aaditya Mattoo, Mushfiq Mobarak, and Nina Yamanis for helpful comments. The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of the World Bank, its Board of Executive Directors, or the governments they represent. The World Bank does not guarantee the accuracy of the data included in this work. The replication package is hosted at <https://data.mendeley.com/datasets/v3sdpfhzj9/1> Please address correspondence to: Gil Shapira, World Bank, Mailstop MC311-311, 1818 H St, NW Washington, DC 20433, USA. E-mail: gshapira@worldbank.org.

This article quantitatively evaluates this trade-off in a susceptible-infected-recovered (SIR) macromodel (e.g., Eichenbaum et al., 2021) augmented with two main features. First, a lockdown can potentially increase child mortality by inducing an economic contraction. The main innovation of our article is to model and quantify this effect. We estimate country-group-specific semielasticities of child mortality with respect to aggregate income changes by applying the methodology of Baird et al. (2011) to microdata from 83 countries, and use the resulting estimates in our quantitative model. Second, we relax the representative agent assumption of most SIR macromodels and allow for three types of agents that differ by age: the children, the working adults, and the elderly (as in Acemoglu et al., 2020).

Infection is assumed to spread through work- and consumption-related activities, as well as community and intrahousehold interactions. Adults are the only ones supplying labor, trading consumption off against the risk to themselves and their family members of contracting COVID-19. A decentralized equilibrium features excess supply of labor since individuals do not internalize the social cost of being infected, which consists of an increased probability of infection for all susceptible individuals as well as a higher infection fatality rate due to limited hospital capacity. A lockdown, which we model as an income tax, reduces labor supply to reduce COVID-19 transmission. A lockdown can lower mortality by either containing the virus or by “flattening the curve,” that is slowing the virus’ spread such that demand for COVID-19 treatment does not exceed health system capacity. However, the reduction in labor supply and consequent consumption losses increases child mortality in low- and middle-income country settings.

We calibrate the model to 85 countries across all income groups. Low-, middle-, and high-income countries differ along several relevant dimensions. First, economic contractions raise child mortality in poorer countries, but not in rich ones. We estimate that a percent decrease in per capita GDP can increase under-5 mortality by up to 0.15 deaths per 1,000 children in the poorest countries. Second, poorer countries have a higher ratio of children to elderly. Since the survival of the former is put at risk by an economic downturn while the latter are most vulnerable to dying from COVID-19, a lockdown in lower-income countries leads to more recession-induced deaths per COVID-19 fatality averted, *ceteris paribus*. Third, a smaller share of social contacts are in the context of work- or consumption-related activities in developing countries compared to developed ones. The preponderance of community-related transmission in low-income countries renders government-mandated lockdowns comparatively less effective at reducing the spread of infections. Finally, low health-care capacity in poorer countries lowers the efficacy of a lockdown through a “flattening the curve” channel as hospitals are quickly overwhelmed.

To highlight the consequences of these country differences, we subject each of the 85 countries in our sample to a uniform reference lockdown that lasts 7 weeks. The duration and the strength of the reference lockdown is chosen based on experiences in seven European countries in the start of the pandemic (Flaxman et al., 2020).¹ We then compare the economic and mortality outcomes in the reference lockdown to a scenario without any government intervention.

Our main quantitative result is that there is substantial variation in health outcomes across countries following the reference lockdown. In our model, the lockdown led to an average of 1.76 child deaths for every COVID-19 fatality averted in low-income countries. The ratio

¹ We take the perspective of the initial months of the COVID-19 pandemic when governments were first confronted with designing pandemic responses in the absence of effective vaccines or therapeutics. We implement only nonpharmaceutical interventions because in the event of a new but similar pandemic, the same trade-offs will be relevant as new vaccines and/or therapeutics would most likely not be immediately available. According to the Oxford Blavatnik COVID-19 Government Response Tracker, every single one of the 175 countries that it covers imposed at least some lockdown policy. Lockdown severity was remarkably similar across countries at different levels of development. On a scale of 0–100, the mean Oxford–Blavatnik lockdown stringency index was 79 in the low-income countries, and 78 in the high-income countries. (These averages are as of April 2020, the most globally synchronized phase of the lockdown. We use the World Bank classification of countries into income groups.)

falls to 0.59 and 0.06 in the case of lower-middle and upper-middle income countries, respectively. By assumption, there is no mortality trade-off in high-income countries. As a result, lockdowns lower the total mortality by 6.2% in the richest countries, but raise total mortality by 2.6% in the poorest ones. The main country characteristics driving the heterogeneity in health outcomes are (i) the semielasticities of child mortality with respect to GDP changes and (ii) demography, as poorer countries are also younger countries.

Finally, we consider a utilitarian approach to designing lockdown policies, in which the social planner trades off COVID-19 deaths averted against loss of life due to reduced GDP and the loss of consumption. The optimal lockdown varies across countries, as marginal costs and benefits are heterogeneous. Poor countries with younger populations generally feature shorter and milder optimal lockdowns, as the governments take into account the impacts on child mortality. Thus, the optimal lockdown significantly reduces the share of the population infected in the rich countries but not in the poorer ones. In the end, the child mortality impact is smaller as well: among the poorest countries, the optimal lockdown would lead to only 0.32 child deaths for every COVID-19 fatality averted.

Our findings hold general lessons for pandemic mitigation, past and future. Empirical evidence has shown repeatedly that in lower- and middle-income countries, infant and child mortality rises more in economic downturns than mortality of other age groups. Thus, the intergenerational trade-off is a generic feature of the policy options in these countries: lockdowns will always adversely affect children disproportionately. The impact of lockdowns on child mortality would then need to be compared to the mortality profile of the disease itself. Evidence suggests that these differ across epidemics. To take some of the most prominent examples, the 1918–1919 influenza pandemic was characterized by an age shift with most excess deaths occurring among young adults (ages 15–44) and fewer excess deaths occurring among those over 65 (Olson et al., 2005; Andreasen et al., 2008). The Ebola case fatality rates (CFRs) for young children under 5 and for elders over 75 are both approximately 80% higher than for prime-age adults (Garske et al., 2017). During the 2003 SARS outbreak, the CFR was estimated to be an order of magnitude higher for patients in China over 60 years of age than those under 40 years (25% vs. 2%) (Jia et al., 2009). A similar age gradient was observed in Hong Kong during the same outbreak (Karlberg et al., 2004). Thus, the exact nature of the intergenerational mortality trade-off will differ across pandemics in a way that is quantifiable within the framework developed here.

Our article complements the burgeoning body of work on the macroeconomic impact of the COVID-19 pandemic (see, among others, Atkeson, 2020; Glover et al., 2020; Alvarez et al., 2021; Barrot et al., 2021; Baqaee and Farhi, 2020, 2021; Bonadio et al., 2021; Kaplan et al., 2020; Krueger et al., 2021). Most closely related are Eichenbaum et al. (2021), who develop an SIR macromodel, and Acemoglu et al. (2020), who model population heterogeneity by age. We customize these macro frameworks to developing countries. Our analysis shares the developing country focus with Alon et al. (2020), Barnett-Howell and Mobarak (2020), Loayza (2020), or Ravallion (2020) who also point out differences between rich and poor countries in the benefits and costs of a lockdown and ultimately come to the same conclusion that the trade-offs are different and country-specific. Our analysis highlights and more importantly quantifies a distinct mechanism, whereby a lockdown potentially increases child mortality in poorer countries. Other work that has surmised the potential toll for infant and child health as a consequence of the COVID-19 pandemic includes Robertson et al. (2020), who use a health care seeking and supply model, and posit reductions in care seeking and available quality of care. In contrast, our approach uses the data on past contractions to calibrate the under-5 mortality semielasticity with respect to the economic shock.

The rest of the article is structured as follows: Section 2 lays out the quantitative framework. Section 3 details the calibration, and 4 presents the quantitative results. Section 5 concludes. The Appendix contains additional details on theory, quantification, and robustness.

2. QUANTITATIVE FRAMEWORK

This section builds a macro-SIR framework along the lines of Eichenbaum et al. (2021) with the added feature that households comprise several members in different age groups (Acemoglu et al., 2020). Our key innovations are (i) to model income shocks as a source of mortality not related to COVID-19; and (ii) to calibrate the model to 85 countries with different underlying characteristics.

2.1. *Economic Environment.* We consider a discrete and infinite time horizon model, $t = 0, 1, \dots, \infty$, and a continuum of households indexed by $j \in [0, 1]$. A model time period corresponds to one week. The measure of households is normalized to 1 in the initial period. Households are formed by individuals differentiated by age group $m \in \{1, 2, 3\}$ to which they belong. Type $m = 1$ individuals are children ages 0–14, type $m = 2$ are working prime-age adults ages 15–59, and $m = 3$ household members are the elderly aged 60 and over. Denote by ℓ^m the mass of individuals of age group m so that $\sum_{m=1}^3 \ell^m = 1$. We omit country indices in the exposition to streamline notation, but the quantitative analysis uses country-specific values for many of the parameters.

Household j evaluates its lifetime utility according to:

$$(1) \quad U_j = \sum_{t=0}^{\infty} \beta^t u(\bar{c}_{jt}, n_{jt}),$$

where β is the discount factor. The instantaneous utility function takes the form

$$(2) \quad u(\bar{c}_{jt}, n_{jt}) = \bar{c}_{jt} - \frac{\theta}{2} n_{jt}^2,$$

where \bar{c}_{jt} is aggregate consumption of household j in period t , and n_{jt} is the amount of labor it supplies.

Household consumption \bar{c}_{jt} aggregates individual consumption $c_{jt}(m)$ of all members of the household:

$$(3) \quad \bar{c}_{jt} = \left[\sum_{m=1}^3 \ell^m (c_{jt}(m))^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

where σ is the elasticity of substitution.

2.1.1. *COVID-19 SIR states.* COVID-19-related health status is relevant for both disease transmission and economic behavior. Each individual can be in one of four states: susceptible (S), infected (I), recovered (R), or deceased (D). One feature of our model is that death can be due to either COVID-19 or another cause. We thus index each household state with integer $k \in \{1, \dots, 64\}$, which uniquely identifies a triplet $\{\zeta(1), \zeta(2), \zeta(3)\}$, where $\zeta(m) \in \{S, I, R, D\}$ indicates the health status of individual m . Appendix Table A.1 reports the list of possible household states.²

2.1.2. *Labor supply, lockdown policy, and government budget.* In our model, only the prime-age adult ($m = 2$) household members supply labor. They are paid a wage w_t , which the

² Note that to reduce the dimensionality of the state space, we make the assumption that within households, all members in the same age group are in the same state. We thus do not allow two children (or adults or elderly) to be in two different states. Although this assumption may be unrealistic at the household level, this simplification does not affect the analysis at the aggregate level.

government can tax at rate μ_t . As in Eichenbaum et al. (2021), the tax rate μ_t will be the instrument by which the policymaker implements a lockdown.³ Tax revenues are then remitted to households in a lump-sum manner. The budget constraint of household j is:

$$(4) \quad c_{jt} = \sum_{m=1}^3 \ell^m c_{jt}(m) \leq \begin{cases} (1 - \mu_t)w_t n_{jt} + \Gamma_{jt}, & \mathbf{S}_{jt}(2) = 1 \text{ or } \mathbf{R}_{jt}(2) = 1 \\ (1 - \mu_t)w_t \phi n_{jt} + \Gamma_{jt}, & \mathbf{I}_{jt}(2) = 1 \\ \Gamma_{jt}, & \mathbf{D}_{jt}(2) = 1, \end{cases}$$

where c_{jt} is total household consumption expenditure. Household income on the right-hand-side of (4) consists of after-tax labor income and the government transfer Γ_{jt} . If the working adult is infected ($\mathbf{I}_{jt}(2) = 1$), the effective labor supply falls by a fraction $\phi \leq 1$. After the death of the working adult ($\mathbf{D}_{jt}(2) = 1$), household j lives off government transfers.

The amount Γ_{jt} transferred to households is determined by the government’s budget constraint, that is,

$$(5) \quad \int_{j=0}^1 \Gamma_{jt} dj = \mu_t w_t \left(\int_{j=0}^1 \mathbf{S}_{jt}(2) n_{jt} dj + \int_{j=0}^1 \mathbf{R}_{jt}(2) n_{jt} dj + \phi \int_{j=0}^1 \mathbf{I}_{jt}(2) n_{jt} dj \right) + \bar{\Gamma}_t,$$

where $\bar{\Gamma}_t$ is some exogenous development assistance revenue.

2.1.3. *Firms.* There is a unit measure of competitive firms that produce consumption goods C_t using the aggregate labor input N_t :

$$(6) \quad C_t = AN_t.$$

Firms choose total labor input to maximize their profit, Π_t :

$$(7) \quad \Pi_t = AN_t - w_t N_t.$$

In equilibrium, goods and labor market clearing conditions are thus

$$\begin{aligned} C_t &= \int_{j=0}^1 c_{jt} dj, \\ N_t &= \int_{j=0}^1 \mathbf{S}_{jt}(2) n_{jt} dj + \int_{j=0}^1 \mathbf{R}_{jt}(2) n_{jt} dj + \phi \int_{j=0}^1 \mathbf{I}_{jt}(2) n_{jt} dj. \end{aligned}$$

2.2. *Mortality and Disease Transmission.* We incorporate a modified SIR model to our macroeconomic framework. In our model, there are three types of mortality risks: (i) economic distress risk, (ii) a COVID-19-related risk, and (iii) an exogenous baseline risk.

2.2.1. *Economic distress and baseline mortalities.* An individual in age group m faces increased mortality during an economic contraction. A contraction is a downward deviation from baseline consumption \tilde{c}_j , defined as the level that would be achieved at time $t = 0$ in the absence of a labor tax ($\mu = 0$). Thus, for household j , a contraction takes place when faced with a positive labor tax or in the case of death of the working adult. In addition, in each period t an individual draws an exogenous age-group-specific baseline death shock with probability $\bar{\pi}_n(m)$.

³ We interpret lockdown policies broadly to include multiple forms of containment, such as encouraging social distancing, restricting public events, school and border closures, and so on. As will become clear below, in the calibration, the lockdown policy will also lower community transmission.

The economic distress mortality cum baseline mortality probability is:

$$(8) \quad \pi_{njt}(m) = \begin{cases} \bar{\pi}_n(m) & c_{jt} \geq \tilde{c}_j \\ \bar{\pi}_n(m) + \nu(m)\left(1 - \frac{c_{jt}}{\tilde{c}_j}\right) & c_{jt} < \tilde{c}_j \end{cases}$$

where $\nu(m)$ is the elasticity of economic distress-related mortality with respect to the fluctuations in consumption. Importantly, in our quantification, $\nu(m)$ will vary by country income level: it will be positive in poorer countries and decreasing in the income level. In rich countries, economic distress-related mortality elasticity $\nu(m)$ will be 0. In the quantification, $\nu(m)$ will be positive only for children ($m = 1$).

2.2.2. *COVID-related mortality.* The infection status of an individual of age group m in household j at time t is denoted as $\mathbf{I}_{jt}(m) \in \{0, 1\}$. Thus, the mass of infected individuals at time t is given by

$$I_t = \int_{j=0}^1 \sum_{m=1}^3 \ell^m \mathbf{I}_{jt}(m) dj.$$

Conditional on contracting COVID-19, the probability of death from the infection takes the following form:

$$(9) \quad \pi_{dt}(m) = \pi_d^m + \kappa^m(I_t),$$

where π_d^m is a baseline infection fatality rate and $\kappa^m(\cdot)$ captures the dependence of mortality on the total infection rate. The function $\kappa^m(\cdot)$ reflects the possibility that a larger epidemic will lead to higher mortality due to saturation of key health services such as ICU beds, oxygen ventilators, etc. (Yang et al., 2020).

In each period, the probability $\pi_{djt}(m)$ that an individual j will die combines both COVID-19 and non-COVID-19 mortality risks. We make the assumption that the economic distress and COVID-19 mortality probabilities are orthogonal to each other in the cross section of households. In that case, the death probability of a person of type m in household j becomes:

$$(10) \quad \pi_{djt}(m) = \begin{cases} 1 - [1 - \pi_{njt}(m)][1 - \pi_{dt}(m)] & \text{if } \mathbf{I}_{jt}(m) = 1 \\ \pi_{njt}(m) & \text{if } \mathbf{I}_{jt}(m) = 0 \end{cases}$$

This equation states the probability of death for an infected individual, thus implicitly conditioning on an applicable household state k (that is, k being one in which household member m is currently infected). To streamline notation, in this equation as well as in (9) and (11), we suppress the conditioning on an applicable k .

2.2.3. *Lockdown policies and COVID-19 disease dynamics.* Adapting the model of Eichenbaum et al. (2021), we assume that the transmission of the infection occurs through four channels: (i) the labor channel, whereby the infection spreads through workplace interactions, (ii) the consumption channel, which comprises contacts occurring while shopping for goods, (iii) the community channel, which represents all other interactions of individuals across households, and (iv) the within-household channel, to account for higher exposure of individuals who share a residence with an infected individual.

A lockdown policy therefore will affect transmission likelihood through these same channels. As a tax on labor income, a lockdown reduces individual labor supply and consequently also household consumption. Both lead to a decrease in infection rates. We further allow a lockdown policy to mitigate community-related transmission with a semielasticity of ξ . This captures restrictions on social gatherings that affect community spread. We do not directly

model decisions related to such gatherings and instead account for the impact of lockdowns on community spread via the parameter ξ .

The probability that a susceptible individual m in household j will get infected in period t is given by:

$$(11) \quad \pi_{I_{jt}}(m) = \begin{cases} \pi_{I1}c_{jt}C_{It} + \pi_{I2}n_{jt}N_{It} + \pi_{I3}I_t(1 - \xi\mu_t) + \pi_{I4}\mathbf{I}_{jt}, & m = 2 \\ \pi_{I3}I_t(1 - \xi\mu_t) + \pi_{I4}\mathbf{I}_{jt}, & m = 1, 3 \end{cases}$$

The first line of (11) describes the infection probability of the working adult. The four terms reflect transmission through consumption, labor supply, the community, and within the household, respectively. Consumption and labor supply transmissions are a function of the aggregate consumption and labor supply of the infected individuals in period t C_{It} and N_{It} , which equal:

$$C_{It} = \int_{j=0}^1 \sum_{m=1}^3 \ell^m c_{jt}(m) \mathbf{I}_{jt}(m) dj,$$

$$N_{It} = \phi \int_{j=0}^1 n_{jt} \mathbf{I}_{jt}(2) dj.$$

Community transmission, on the other hand, is a function of the total number of infected people, I_t , as defined above. Finally, for within-household transmission, \mathbf{I}_{jt} , equals to 1 if any member of the household j is infected, and 0 otherwise. The second line of Equation (11) applies to the children and the elderly, who will only be infected through the community or within-household transmission channels, since they do not work and are assumed not to get exposed through consumption-related activities.

The total number of newly infected individuals T_t is thus given by:

$$(12) \quad T_t = \int_{j=0}^1 \sum_{m=1}^3 \ell^m \pi_{I_{jt}}(m) \mathbf{S}_{jt}(m) dj,$$

where $\mathbf{S}_{jt}(m)$ is an indicator function that takes the value of 1 when member m of household j is susceptible, and 0 otherwise. The number of susceptible individuals, S_t , evolves according to

$$(13) \quad S_{t+1} = S_t - T_t.$$

In period t , all infected individuals will receive the “recovery” shock. With probability $\pi_{rt}(m)$, the member recovers, with probability $\pi_{dt}(m)$ she/he dies, and with probability $\pi_{it}(m)$, she/he stays infected. Note that $\pi_{rt}(m) + \pi_{dt}(m) + \pi_{it}(m) = 1$. The number of infected individuals thus evolves according to:

$$(14) \quad I_{t+1} = \int_{j=0}^1 \sum_{m=1}^3 \ell^m \mathbf{I}_j(m) \pi_{it}(m) dj + T_t,$$

which consists of previously infected people who remain so for one additional period and newly infected individuals.

2.3. Household Optimization. We now turn to household optimization, subject to the aggregate state of the economy as summarized by $\Theta_t = \{C_{It}, N_{It}, I_t\}$ and government policy μ_t . We first note that all the households in state k face the same maximization problem and make the same decision. As a result, we use the subscript k instead of j to indicate the variables for a household in state k .

2.3.1. *Consumption and labor supply.* Before solving the dynamic problem, we first solve the within-period problem through backward induction by expressing household instantaneous utility as only a function of labor supply and then optimizing accordingly.

As we abstract from saving and risk sharing across households, the solution to the consumption problem is static and is characterized by the binding budget constraint (4) for household k . In light of this observation, we can rewrite the probability of death specified in Equation (10), $\pi_{dkt}(m)$, as a function of labor supply:

$$(15) \quad \pi_{dkt}(n_{kt}; m) = \begin{cases} \bar{\pi}_n(m) + v(m) \left[1 - \frac{c_{kt}(n_{kt})}{c_k} \right] + \pi_d^m + \kappa^m(I_t) & \text{if } \mathbf{I}_{kt}(m) = 1 \\ \bar{\pi}_n(m) + v(m) \left[1 - \frac{c_{kt}(n_{kt})}{c_k} \right] & \text{if } \mathbf{I}_{kt}(m) = 0 \end{cases}$$

where c_{kt} depends on n_{kt} through (4), and is a function of the infection state of the working adult. Similarly, the probability of a susceptible adult contracting COVID-19 can also be reexpressed as a function of labor supply:

$$(16) \quad \pi_{Ikt}(n_{kt}; 2) = \pi_{I1}[(1 - \mu_t)w_t n_{kt} + \Gamma_{kt}]C_{It} + \pi_{I2}n_{kt}N_{It} + \pi_{I3}I_t + \pi_{I4}\mathbf{I}_k.$$

Finally, the standard property of CES aggregation implies that we can also rewrite the flow utility function as a function of labor supply:

$$u(\bar{c}_{kt}, n_{kt}) = \left[\sum_{m=1}^3 \ell^m [c_{kt}(m)]^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} - \frac{\theta}{2} n_{kt}^2,$$

in which the individual-level consumption is equal to

$$c_{kt}(n_{kt}; m) = \frac{\mathbf{1}_{kt}(m)}{\sum_{m'=1}^3 \ell^{m'} \mathbf{1}_{kt}(m')} c_{kt}(n_{kt}),$$

and where once again the dependence of c_{kt} on n_{kt} takes the form of (4). Combining the two and simplifying:

$$(17) \quad u(n_{kt}) = c_{kt}(n_{kt}) \left[\sum_{m=1}^3 \ell^m \mathbf{1}_{kt}(m) \right]^{\frac{1}{\sigma-1}} - \frac{\theta}{2} n_{kt}^2.$$

2.3.2. *Dynamic optimization.* With the solution of the consumption problem in hand, we can turn to the dynamic problem of a household in state k . The Bellman equation for household in state k can be written

$$(18) \quad U_{kt}(\Theta_t) = \max_n u(n) + \beta \sum_{k'=1}^K \rho_{kk'}(n|\Theta_t) U_{k',t+1}(\Theta_{t+1}),$$

subject to the transition probabilities from state k to k' , $\rho_{kk'}(\cdot)$, which depend on the aggregate state of the economy Θ_t .

The first-order condition determines optimal labor supply n_{kt} :

$$(19) \quad u'(n_{kt}) + \beta \sum_{k'=1}^K \frac{\partial}{\partial n} \rho_{kk'}(n_{kt}|\Theta_t) U_{k',t+1}(\Theta_{t+1}) = 0.$$

We can then write the optimal labor decision as

$$(20) \quad n_{kt} = \frac{\left[\sum_{m=1}^3 \ell^m \mathbf{1}_{kt}(m) \right]^{\frac{1}{\sigma-1}} (1 - \mu_t) w_t + \Lambda_{kt}}{\theta},$$

where $\Lambda_{kt} = \beta \sum_{k'=1}^K \frac{\partial}{\partial n} \rho_{kk'}(n_{kt} | \Theta_t) U_{k',t+1}(\Theta_{t+1})$.

First, note that in the absence of capital accumulation, households' labor decisions will only affect a subset of the transition probabilities, $\rho_{kk'}$. Labor supply only affects the probability of infection of the working age adults and the non-COVID mortality rate of children (since we assume that excess mortality from an economic contraction only affects children). All the other mortality rates, infection rates, and recovery rates follow a process that is not influenced by the decision of atomistic agents but depend on the aggregate state Θ_t . We thus state the following result (proof in the Appendix):

LEMMA 1 (FIRST-ORDER CONDITIONS). *The first-order condition for the household in state k , period t is:*

$$(21) \quad \frac{\partial u}{\partial n_{kt}} - \lambda_{\pi_t}^k [\pi_{I1}(1 - \mu_t) w_t C_{It} + \pi_{I2} N_{It}] + \lambda_{\pi_d}^k v(1) \frac{(1 - \mu_t) w_t}{\tilde{c}_k} = 0,$$

where $\lambda_{\pi_t}^k$ is the Lagrangian multiplier on the infection probability, $\pi_{Ikt}(2)$:

$$(22)_{\pi_t}^k = \beta(1 - \pi_{nkt}(2)) \left\{ \sum_{k'(2)=S} \rho_{kk'}(1 | \Theta_t) \rho_{kk'}(3 | \Theta_t) U_{k'} - \sum_{k'(2)=I} \rho_{kk'}(1 | \Theta_t) \rho_{kk'}(3 | \Theta_t) U_{k'} \right\} \geq 0,$$

and $\lambda_{\pi_d}^k$ is the Lagrangian multiplier for the non-COVID child mortality rate, $\pi_{dkt}(1)$:

$$(23) \quad \lambda_{\pi_d}^k = \beta \left\{ \sum_{k'(1) \neq D} \rho_{kk'}(2 | \Theta_t) \rho_{kk'}(3 | \Theta_t) U_{k'} - \sum_{k'(1)=D} \rho_{kk'}(2 | \Theta_t) \rho_{kk'}(3 | \Theta_t) U_{k'} \right\} \geq 0.$$

For households without a susceptible prime-age adult, $\lambda_{\pi_t}^k = 0$ as the terms inside the curly bracket in Equation (22) are equal to zero. Similarly, $\lambda_{\pi_d}^k = 0$ for households without a child as the terms inside the curly bracket in Equation (23) are equal to zero.

The first-order condition captures the trade-off between the static optimization (i.e., today's consumption vs. leisure) and the health risk of increased exposure through consumption and work. Lemma 1 describes the heterogeneity in households' responses to the pandemic, as a function of their demographic composition and the health status of their members. On the one hand, labor decisions have no dynamic implications when no prime-age adult is susceptible. On the other hand, incentives to increase (decrease) labor supply depend on whether there are children (susceptible elderly people) in the household.

2.3.3. *Equilibrium.* An equilibrium of the economy in period t is defined by a vector of labor supply decisions $\{n_{kt}\}_{k \in \{1, \dots, K\}}$ such that n_{kt} is a solution to (18) for some given Θ_{t+1} , and Θ_{t+1} is, in turn, determined by transition probabilities (10) and (11). To solve for the equilibrium, therefore, we propose the following algorithm:

2.3.4. *Solution algorithm.* Take the policy vector, μ_t , as given. Start with a guess of n_{kt} for all $k = 1, \dots, K$ and $t = 1, \dots, \mathcal{T}$.

1. Given the initial conditions, simulate the model forward from $t = 1$ to \mathcal{T} to generate S_t , I_t , R_t , and D_t , as well as all the transition probabilities.
2. Infer U_{kt} for every k via the following steps. The details are discussed in the Appendix.
 - (a) Compute the postpandemic steady-state values of U_k for all k .
 - (b) Compute backward from the postpandemic state \mathcal{T} to 1 for all the U_{kt} .
3. Infer $\lambda_{\pi_d^k, t}$ and $\lambda_{\pi_I^k, t}$ from the first-order conditions of $\pi_{dkl}(1)$ and $\pi_{Ikl}(2)$, conditional on U_{kt} .
4. Infer n_{kt} from the first-order conditions of n_{kt} . Iterate on n_{kt} until convergence.

2.4. *Discussion.* Before moving on to the calibration and quantification, we discuss some features of our theoretical framework.

2.4.1. *Non-COVID-19 mortality.* In our framework, economic downturns only affect the mortality of children, and only in the poorer countries. This combination of assumptions appears the most realistic in light of available empirical evidence. Mortality patterns in low- and middle-income countries are typically found to be counter-cyclical, especially for vulnerable groups such as infants and young children (see, e.g., Ferreira and Schady, 2009, for a review of the available literature). On the other hand, there is a broad consensus that adult mortality in high-income countries is if anything procyclical (e.g., Ruhm, 2000; Stevens et al., 2015). Unfortunately, there is very little existing evidence on the cyclicity of mortality in nonchild populations in low- and middle-income countries. Allowing for a mortality increase among nonchildren due to lockdowns would only increase the ratio between the non-COVID-19 deaths caused by the lockdowns and the number of COVID-19 deaths averted, conditional on a given lockdown severity. However, we would expect the overall impact via the elderly mortality rates to be much smaller. One reason is that, as we highlight below, adults above 60 account for a small proportion of the population in developing countries.

2.4.2. *Consumption.* In our framework, the child's consumption depends only on the total income of the household. It could be that in an economic downturn, the consumption of young children falls by even more than average household consumption for various reasons such as the household privileging the consumption of the main earner. Unfortunately, we cannot incorporate this possibility directly because there is no sufficiently reliable data to inform within-household consumption differences over the business cycle. If one believes that children's consumption falls more than proportionally to household consumption during economic downturns, our approach is conservative and allowing for this possibility would only quantitatively strengthen our main point.

2.4.3. *Health system oversaturation.* Our framework models the pandemic's impact on child mortality through the negative income shock, which can include changes in the utilization of health services as experienced historically during nonpandemic economic downturns. However, the pandemic may impact non-COVID-19 child mortality through changes in availability or utilization of health services in unique ways. Indeed, if health system oversaturation due to treating COVID-19 patients reduces the coverage of life-saving services for children, then ceteris paribus child mortality would increase in the absence of a lockdown. This would reduce the ratio between excess child deaths and COVID-19 mortality averted by a given lockdown policy. On the other hand, there is suggestive evidence that lockdowns themselves might also reduce utilization of essential health services. In a study of 18 low- and middle-income countries, Ahmed et al. (2021) find a correlation between monthly reductions in health service utilization and stringency of mobility restrictions, even controlling for the monthly COVID-19 burden. If lockdowns reduce coverage of life-saving health services, then we would be underestimating the ratio between excess child deaths and COVID-19 mortality averted by a given lockdown. As for reduced demand for services out of fear of

contracting the disease in health facilities, it may happen independent of whether lockdowns are imposed. In sum, it is not clear whether non-COVID-19 mortality due to reduced health service utilization (or reduced service quality) would be higher with lockdowns or with health system congestion.

2.4.4. State capacity. In our framework, the lockdown is successfully implemented by the government. One may be worried that if a particular country does not have the state capacity to enforce a lockdown, then our quantification is of limited policy relevance to that country. This is a general critique of policy analysis in low-state-capacity environments. Nevertheless, emerging empirical evidence suggests that indeed, lockdowns reduced economic activity in even the lowest income countries (Aminjonov et al., 2021; Beyer et al., 2021). In addition, throughout our analysis, there continues to be community transmission that is not entirely eliminated by the lockdown, and we calibrate community transmission to be higher in developing countries. This feature of our quantification reflects among other things lower state capacity in poorer countries.

2.4.5. Government transfers and the wealth effect on labor supply. In our model, the government rebates the lockdown-tax income back to households lump sum. This assumption can be thought of as a stand-in for the transfer programs that governments put in place jointly with the lockdowns. In principle, there may be a wealth effect from these government transfers on labor supply. If the wealth effect is large, the labor supply response to the lockdown could be sensitive to the adopted assumption on government transfers. In our model, the GHH (Greenwood et al., 1988) functional form for preferences rules out a substantial wealth effect on labor supply, and thus, the shape of the labor supply response to the lockdown is not sensitive to the assumption we put on the transfers. This is consistent with the available empirical evidence that there is little to no change in adult labor supply as a result of cash transfer program receipt in low- and middle-income countries (see Baird et al., 2018, for a review of the literature).

3. DATA AND CALIBRATION

The strength and duration of a lockdown are critical aspects of our quantitative analysis. Our reference lockdown policy attempts to mimic what had been observed in the first weeks of the pandemic. It is henceforth defined by three parameters: its starting time, length, and severity. To calibrate these parameters alongside transmission rate parameters, we proceed in two steps. In the first step, we calibrate the transmission parameters $\pi_{I1}, \dots, \pi_{I4}$ to match the relative importance of the different transmission modes and the overall predicted infection rate in an unmitigated spread scenario. In the second step, we calibrate the effect of lockdown severity on community transmission ξ , alongside the strength of the reference lockdown $\bar{\mu}$, to jointly match the decline in GDP and the reproduction number R_0 as estimated by Flaxman et al. (2020) for European countries early in the pandemic.

3.1. Infection and Mortality Parameter Calibration. The within-household transmission parameter, π_{I4} , is taken from a meta-analysis of household transmission estimates from different settings. Lei et al. (2020) estimate the secondary infection rate in the household to be 0.27.

To discipline the three other transmission parameters, we jointly match three moments. The first moment is the proportion of the population that would get infected in each country in the absence of any mitigating policy. We use projections reported by Walker et al. (2020) that use country- and age-specific contact patterns to simulate health impacts of COVID-19 in 202 countries. They develop an SIR model incorporating the age distribution of each country. Employing a basic reproduction number (R_0) of 3.0, they project that about 90% of the population would ultimately either recover from infection or die in an unmitigated epidemic scenario in lower- and middle-income countries. The unmitigated epidemiological model in Walker

et al. (2020) assumes no behavioral response to the pandemic. For consistency, we assume that households continue to supply labor and consume at the same levels as in the prepandemic steady state in this stage of the calibration exercise. This assumption is relaxed in the following steps of the calibration.

The other two moments used for the calibration of the transmission parameters are the shares of infections occurring through work- and consumption-related activities. As most of the world's population lives in urban areas, we chose to use data reported by Johnstone-Robertson et al. (2011) on locations of contacts in a South African township community. The authors define close contacts as those involving physical contact or a two-way conversation with three or more words. They find that 6.2% of close contacts occur in workplaces, whereas 3.5% occur in shops or local bars and therefore can be thought of as related to consumption. Another 8.9% of close contacts take place during transport and could theoretically be linked to either labor or consumption. We assume that half of the transport contacts is related to labor and half to consumption. This implies that 10.6% of close contacts are related to work and 8% to consumption. For high-income countries, we use the rates employed by Eichenbaum et al. (2021) for the United States. Based on an analysis of the Bureau of Labor and Statistics 2018 Time Use Survey data and contact patterns reported in Ferguson et al. (2006) and Lee et al. (2010), they conclude that 16% of transmissions are related to consumption and 17% to work.

3.2. Reference Lockdown and Community Transmission Parameter. Conditional on the transmission parameters calibrated above, we calibrate the reference lockdown policy, and the parameter determining the relationship between lockdown strength and reduction in community transmission, ξ .

A country starts to impose the reference lockdown when the infected population reaches 2.6%. This rate is based on the COVID-19 prevalence at the time of the first lockdown in the Italian municipality of Vo, the site of the first COVID-19-related death detected in Italy (Lavezzo et al., 2020). In our calibration, the countries start to impose the reference lockdown policy between weeks 9 and 13, with an average start date at week 11.

The length of the reference lockdown policy is based on Flaxman et al. (2020) that estimates the impacts of nonpharmaceutical interventions in 11 European countries during the first months of the pandemic.⁴ We drop the four countries that only imposed mild or no lockdown policies (Denmark, Norway, Sweden, and Switzerland) and work with the remaining seven countries.⁵ We compute the lockdown length for each country based on the difference between the reported lockdown date and the end of the sample period in Flaxman et al. (2020). The lockdown policies range between 43 and 54 days, with an average of seven weeks, which we use as the length of the reference lockdown. Appendix A.2.1 shows that our results are robust to variations in length and the starting time of the reference lockdown.

The strength of the lockdown was inferred from the GDP decline in the first two quarters of the year 2020. As explained later in this section, we calibrate $\theta = 1$, which implies that a lockdown policy of $\bar{\mu}$ reduces aggregate labor supply and GDP by $\bar{\mu}$. Therefore, an x -day lockdown reduces the two-quarter GDP by $x\bar{\mu}/180$, from which one can infer $\bar{\mu}$, conditional on the length of the lockdown calibrated above and the observed decline in GDP. For example, the inferred $\bar{\mu}$ in Germany with a 6.68% decline in GDP and a 43-day lockdown is $0.0668 \cdot 180 / 43 = 0.2796$. We repeat the calculation for all seven countries and find that the

⁴ Out of the five nonpharmaceutical interventions studied in Flaxman et al. (2020), we focus on “lockdowns” to quantify the reference lockdown. Note that the “lockdown” in our model should be interpreted broadly to include the other four forms of intervention policies in Flaxman et al. (2020), such as social distancing, self-isolation, school closure, and restricting public events.

⁵ The four countries with mild or no lockdown policies lead to corner solutions in the calibration of ξ . If the country imposes a mild lockdown, the implied $\bar{\mu}$ for this country will be low, which makes it impossible to reach the target R_0 even with $\xi = \infty$. The remaining seven countries are Austria, Belgium, Britain, France, Germany, Italy, and Spain.

strength of the lockdown policy to be between 28% and 46%. The average across the seven countries is 38%, which we use as $\bar{\mu}$ for the reference policy.

Given the country-specific $\bar{\mu}$ and lockdown length in these seven countries, we can then compute ξ , the elasticity of community transmission with respect to lockdown, for each of these countries. To do this, we simulate the model and target the postlockdown R_0 of 0.66 as reported in Flaxman et al. (2020) for the sample countries. The corresponding R_0 at period t in our model is computed as $\frac{I_{t-1}/I_{t-1}}{\pi_{rt} + \pi_{dt}}$, where π_{rt} and π_{dt} are the population-weighted average recovery and mortality rates in period t . We take the R_0 at the period after the lockdown policy ends, as the counterpart of the postlockdown reproduction number in Flaxman et al. (2020).⁶ The resulting ξ ranges between 1.9 and 3.5 among the seven countries with an interior solution, and we take the average of 2.32 to apply to all the 85 countries in the full sample.

3.3. *Mortality Rates.* We look to two distinct literatures to calibrate our mortality parameters.

3.3.1. *COVID-19 mortality.* Walker et al. (2020) project hospitalization and mortality rates per age group that are, in turn, based on findings from China reported by Verity et al. (2020).⁷ Conditional on infection, the average projected hospitalization and mortality rates in low- and middle-income countries are listed in Appendix Table A.2. In the calibration of the model, we use country-specific rates as countries have different age distributions within the broader age groups defined in our model.

For severe cases of COVID-19 infection, hospitalization offers treatments such as oxygen therapy for patients with respiratory failure. Therefore, it is believed that when hospital care cannot be accessed, the CFR for a COVID-19 infection is higher. We assume that the COVID-related mortality is elevated by a factor of 3 for those patients who are in need of hospitalization but cannot receive it.⁸

We denote by $\pi_b(m)$ the probability an individual of group m requires hospitalization, conditional on being infected. The share of individuals in need of hospital beds at time t , B_t , is given by

$$B_t = \pi_b(1)I_t(1) + \pi_b(2)I_t(2) + \pi_b(3)I_t(3).$$

⁶ We have computed the R_0 for all the countries in the sample using our calibrated model. The postlockdown R_0 is similar across countries and is close to those reported for the European countries in Flaxman et al. (2020). The average R_0 among the 85 countries one week after the reference lockdown is 0.57. In comparison, the same statistic among the 11 European countries in Flaxman et al. (2020) is 0.66. Given that the R_0 estimate in Flaxman et al. (2020) is our calibration target, the consistency among 85 countries provides some reassurance, but not external validation. Our estimates are also broadly consistent with the R_0 's reported in the economics literature, such as Atkeson et al. (2020) and Fernández-Villaverde and Jones (2020). Although our estimated R_0 is not directly comparable to the R_0 reported in the papers above—the main difference is that the R_0 's in our article are computed based on simulated data assuming a seven-week reference lockdown, while the R_0 's reported in these two papers are based on the actual data where the strength and length of the lockdown policies vary—the general consistency is reassuring.

⁷ Infection fatality rates calculated with data from China might not be generalizable to other countries because of factors such as prevalence of comorbidities and quality of health services. Meta-analyses found mortality rates in line with those reported by Verity et al. (2020) (Levin et al., 2020; Meyerowitz-Katz and Merone, 2020). It is important to note, however, that these analyses overwhelmingly rely on studies from high-income countries.

⁸ For example, Yang et al. (2020) find that the CFR by the middle of February 2020 within the city of Wuhan, China, the presumed outbreak location of origin for COVID-19, was substantially higher (5.25%) than in regions outside Wuhan but within the same province (1.41%) and in regions outside the province (0.15%). Although the CFRs within and outside Wuhan should vary for several reasons, including availability of testing, a key factor was the initial demand for hospital beds exceeded the supply.

We assume that hospital bed allocation is random among those in need. Denoting by h the number of hospital beds, the probability that an infected individual dies at period t is given by

$$\pi_{idt}(m) = \begin{cases} \pi_{ni}(m) + \pi_d(m) & \text{if } B_t \leq h \\ \pi_{ni}(m) + (1 - \pi_b(m))\pi_d(m) + \pi_b(m)\left(\frac{k}{B_t}\pi_d(m) + 3\frac{B_t-k}{B_t}\pi_d(m)\right) & \text{if } B_t > h \end{cases}$$

As in Atkeson (2020) and Eichenbaum et al. (2021), we assume that it takes an average of 18 days from infection to either recover or die. To obtain weekly mortality probabilities, we multiply the rates obtained from Walker et al. (2020) by 7/18.⁹

The number of hospital beds in each country is obtained from the World Bank’s World Development Indicators. It should be noted that the indicator is not measured frequently, particularly in lower-income countries. We use the most recent measurement reported for each country. In our sample, the average number of hospital beds per 1,000 people is 0.6, 1.6, 3, and 4.15 in low-, lower-middle, upper-middle, and high-income countries, respectively.

3.3.2. Non-COVID-19 mortality. Baseline mortality rates $\bar{\pi}_{n1}$, $\bar{\pi}_{n2}$, and $\bar{\pi}_{n3}$ are computed from country-specific life table data obtained from the Global Health Observatory Data Repository of the World Health Organization.

In terms of elevated mortality due to shortfalls in aggregate income, several papers have estimated the relation between economic shocks and infant or young child mortality (Bhalotra, 2010; Baird et al., 2011; Cruces et al., 2012; Friedman and Schady, 2013). For low- and middle-income countries, the population groups most vulnerable to declines in aggregate income are young children and, perhaps, the elderly (Cutler et al., 2002). We focus on mortality impacts among children under-5 as this population group has been the most extensively studied. We estimate the effect of short-term aggregate income shocks on mortality following the methodology of Baird et al. (2011). We use data on GDP per capita from the World Development Indicators. The values are adjusted for purchasing power parity, corresponding to 2011 U.S. dollars. Data on infant and child mortality are taken from retrospective birth histories as reported in the Demographic and Health Surveys (DHS) conducted in 83 low- and middle-income countries between 1985 and 2017. The combined sample is of 5.2 million births in low- and middle-income countries. We run regressions of the following form:

$$D_{ict} = \alpha_c + \beta \log \text{GDP}_{ct} + f_c(t) + \delta_c + \epsilon_{ict},$$

where D_{ict} is a binary indicator that takes the value 1 if child i in country c died in year t , $\log \text{GDP}$ is the natural logarithm of per capita GDP, $f_c(t)$ is a country-specific flexible time trend, δ_c is the country fixed effect, and ϵ_{ict} is the error term. Standard errors are clustered at the country level.

We run the regression separately for countries of different income levels, as classified by the World Bank 2020 income groups. The main result is that a 1% decrease in per capita GDP is associated with a 0.15 increase in under-5 mortality per 1,000 children in low-income countries. The semielasticity is 0.10 and 0.03 for lower- and upper-middle-income countries, respectively. We assume that under-5 mortality is not impacted by income shocks in high-income countries. Unlike the results from low- and middle-income countries, studies analyzing data from the United States find mortality to be procyclical (Ruhm, 2000; Dehejia and Lleras-Muney, 2004).

⁹ It has been noted that the Eichenbaum et al. (2021) SIR macromodel misses on the timing of COVID-19 deaths for the United States. Unfortunately, there is no practical way to assess whether our model reproduces the relative timings of the economic downturn versus COVID-19 deaths in our sample of countries. The first basic issue is that our exercise does not model actual lockdowns that occurred in these 85 countries. Second, a systematic comparison of our model’s predictions to actual lockdown outcomes would be complicated by the large heterogeneity among countries in the timing of the appearance of the virus, seasonality in transmission due to differences in climate, and the timing of the policy responses.

To map the estimated semielasticities into our calibration, we define s_{0-5} to be the share of children under five years old in the total number of children of ages 0–15. The semielasticity of child mortality with respect to consumption is given by

$$v_g(1) = s_{0-5}\beta_g, \quad g = \text{LIC, LMIC, UMIC},$$

where β_g represents the regression coefficients for low-, lower-middle-, and upper-middle-income countries. $v(1)$ equals zero in high-income countries.

It is important to note that we convert the annual semielasticity estimate into weekly frequency in the quantification, to match the period definition in the model. Unfortunately, we are unable to estimate the mortality semielasticity at a shorter frequency due to data limitations—the national accounts data for many developing countries are only available annually. However, we believe that the underlying relationship operates at a higher frequency than an annual frequency for several reasons. Baird et al. (2011) find that only contemporaneous GDP deviations are correlated with mortality likelihood even though a large share of the infants in the estimating data set experienced the majority of the in utero period in the lagged year and an equal share of infants experienced the majority of their first year of life in the leading year. Moreover, the authors find that the coefficients on economic conditions in utero and after the first month of life are both small and insignificant. By contrast, the coefficient on per capita GDP in the first month is large, significant, and very close in magnitude to the main effect reported in the article. These results underscore that it is the economic conditions around birth (say the last months of pregnancy and the first months of life) that matter most for infant survival during economic contractions. As 47% of under-5 mortality in 2019 occurred in the neonatal period, it appears clearly that the actual frequency through which economic contraction affects child survival is shorter than an annual frequency.

The empirical estimate of the mortality semielasticity with respect to GDP downturns will be mapped in the model to the increased mortality from lower consumption. However, the estimated coefficient may capture other mechanisms, such as lower availability or utilization of health services. Unfortunately, there is insufficient health system capacity time series data to isolate this effect empirically. The available health capacity data would be de facto absorbed by the country fixed effects in estimation. Going further, we tested for heterogeneous semielasticities when stratifying by hospital capacity and found no statistically significant differences in the semielasticity estimates across countries with different levels of hospital capacity.

Further, we believe it unlikely that the health capacity channel is an important one for the magnitude of the semielasticity estimate. This is because the semielasticity is estimated on past (prepandemic) data on “regular” economic downturns, instead of those brought on by pandemics. In a regular economic downturn, there is less reason to believe that the recession in and of itself leads the health system capacity to be overburdened. To the extent that lower utilization of health services is responsible for part of the child mortality response in regular recessions, it is likely due to lower demand for health services by the households suffering negative income shocks, instead of sharp reductions in availability.

3.4. Demographic and Economic Parameters. Country-specific age distributions are obtained from the 2020 World Population Prospects. The age distribution is used to compute s_{0-5} , which is then used to rescale the semielasticity of under-5 mortality to the age group 0–15 using the formula in the previous section. In addition, we use the age distribution of the three age groups to compute the masses of the different age groups within the household (ℓ^1, ℓ^2, ℓ^3) in each country.

The weekly discount factor equals to $\beta = 0.96^{1/52}$ to reflect an annual risk-free rate of 4%. We assume that at $t = 0$, $\varepsilon = 0.1\%$ of population is infected. We set $\phi = 0.8$ so that an infected working adult is only 80% effective in supplying labor. This is equivalent to assuming that 80% of the infected prime-age population is either asymptomatic or experiences a mild

case.¹⁰ We set θ to 1 so that the steady-state labor supply in the prepandemic world is normalized to 1 in all countries.

The parameter σ governs the elasticity of substitution between household members. We set it to 3, so that the loss of a nonproductive household member—the children or the elderly—with mass ℓ^m reduces the instantaneous utility, $u(\cdot)$, by a proportion ℓ^m in steady state. Appendix A.1.3 provides the details of the derivation. Appendix A.2.2 discusses alternative strategies to calibrate this parameter and shows that our results are robust to the variations in σ .

The values of $\bar{\Gamma}_t$ are calibrated based on the Atlas of Social Protection Indicators of Resilience and Equity (ASPIRE) data set. In the data set, 1.55% of GDP was spent on social assistance programs on average. We assume that $\bar{\Gamma}_t$ is constant across the entire simulation, and calibrate it to be 1.55% of GDP in every country.

Table 1 summarizes all the parameters of the model and indicates the data sources used to calibrate them. Appendix Table A.3 lists the countries in the sample. Appendix Tables A.4 and A.5 list all the country-specific parameters for each country.

4. RESULTS

To quantitatively illustrate how the same policy might lead to different mortality outcomes across countries, we compare two scenarios. The first scenario traces economic and disease-related behavior without any government intervention. The second scenario involves the reference lockdown as described above, where a labor tax of 38% is imposed for a seven-week period once the rate of infection prevalence reaches 2.6%. Although the reference lockdown is picked to mimic policies adopted during the early months of the pandemic, it is not designed to capture all the complexities of mobility and social gathering restrictions imposed by various countries. Rather, the results below aim at highlighting the large heterogeneity in outcomes following the adoption of the same policy rule.

4.1. Lockdowns and Total Mortality. Figure 1 plots the reduction in adult COVID-19-related mortality as a result of the reference government-imposed lockdown. The figure depicts excess adult mortality in the first year of the pandemic under the reference lockdown, relative to excess adult mortality in the no-action scenario. In both pandemic scenarios, excess adult mortality is the difference between the number of adult deaths and the number of adult deaths had the economy not experienced a COVID-19 outbreak. Overall, a single seven-week lockdown will reduce adult mortality from COVID-19 by less than 9% in all countries. The figure also shows that the efficacy of the reference lockdown at averting mortality is correlated with per capita income. In low-income countries, an average of 3.5% of COVID-19-related deaths are averted, in comparison to an average of 6.2% in high-income countries.¹¹ Several factors drive this pattern. First, wealthier countries' populations have a larger share of adults over 60, the group most at risk of dying from COVID-19. Second, because of greater hospital capacity in wealthier countries, a slowed pace of the virus' spread is more likely to translate into higher survival probabilities. Finally, greater shares of transmission in high-income countries occur through labor and consumption-related contacts. Therefore, the

¹⁰ This corresponds to media reports on the results of an unpublished seroprevalence study conducted in New York in April 2020 ([nytimes.com/2020/04/23/nyregion/coronavirus-antibodies-test-ny.html](https://www.nytimes.com/2020/04/23/nyregion/coronavirus-antibodies-test-ny.html)). Li et al. (2020) and Stringhini et al. (2020) report even higher rates of asymptomatic cases in China and Switzerland.

¹¹ The lockdown slows the virus' spread but over a full-year horizon, the effect on the total share of the population ever infected is small in many countries. On average, the number of cases averted by the reference lockdown is around 2% of the country's total population. Appendix Figure A.1 plots the reduction in total infections as a fraction of population against per capita income. The reference lockdown is more effective at averting cases in the high-income countries: it prevents 2.8% of the population in high-income countries from getting infected, but only 1.8% in low-income countries.


TABLE 1
PARAMETERS

Name	Value	Source/Target	Note
π_{I1}	Country-specific	Share of transmission due to consumption-related activities	Consumption-related activity transmission
π_{I2}	Country-specific	Share of infection due to work-related activities	Work-related transmission
π_{I3}	Country-specific	Proportion of infected population	Community-based infection
π_{I4}	0.27	Lei et al. (2020)	Within-household infection
ξ	2.32	post-lockdown R_0 in Flaxman et al. (2020), seven-country sample	Impact of lockdown on Community transmission
–	Country-specific	population infection rate of 2.6%	Start date of the reference lockdown
–	7	lockdown length from Flaxman et al. (2020), seven-country sample	Length of the reference lockdown
$\bar{\mu}$	38%	GDP decline in 2020Q1-Q2 in the seven countries from Flaxman et al. (2020)	Strength of the reference lockdown
semielasticity of child mortality w.r.t. income			
$v_j(1)$	0.15	Authors' estimation using DHS data	Low-income countries
	0.10	Authors' estimation using DHS data	Lower-middle-income countries
	0.03	Authors' estimation using DHS data	Upper-middle-income countries
	0.00	Assumed based on literature	High-income countries
$\pi_d(m)$	Country-specific	Walker et al. (2020)	COVID-19 mortality, type- m
$\pi_h(m)$	Country-specific	Walker et al. (2020)	Probability of hospitalization, type- m
k	Country-specific	World Development Indicators	number of hospital beds per 1000 people
$\bar{\pi}_n(m)$	Country-specific	WHO Global Health Repository	Non-COVID mortality, type- m
$\ell(m)$	Country-specific	World Population Prospects, 2020	Mass of type- m
β	$0.96^{1/52}$	annual risk-free rate of 4%	Weekly discount factor
ε	0.1%	Eichenbaum et al. (2021)	Initial infected population
ϕ	0.8	Eichenbaum et al. (2021)	Efficiency loss due to infection
θ	1.0	steady-state labor supply = 1	Disutility from labor supply
σ	3.0	loss of type- m reduces $u(\cdot)$ by ℓ^m	Elasticity of substitution between household members
$\bar{\Gamma}_t$	1.55% of GDP	ASPIRE data set	External transfer payment

NOTE: This table lists the calibrated parameters discussed in the main text. Country-specific parameters are presented in the Appendix.

reduced economic activity in these countries has a bigger impact on the virus transmission relative to countries where a larger share of transmissions occurs through community contacts.


4.1.1. *Lockdowns and the mortality trade-off.* Panel (a) of Figure 2 illustrates the intergenerational mortality trade-off that is the focus of this article. As in Figure 1, on the horizontal axis is log per capita income. The vertical axis represents the number of children's lives lost during the first year of the pandemic per COVID-19 fatality averted by the reference lockdown. There is a pronounced negative relationship between this indicator and income. By construction, no child life is lost due to COVID-19-related lockdowns in high-income countries, where we assume that GDP contractions have no impact on child mortality. High-income countries therefore lie on the horizontal axis. For lower-income countries, however, there can be a substantial loss of children's lives for each averted COVID-19 fatality. In 19 of the low- and lower-middle income countries in our sample, the reference lockdown policy



NOTES: This figure displays the ratio of COVID-19 fatalities with and without the reference lockdown against the logarithm of PPP-adjusted per capita GDP. On the vertical axis is the adult COVID-19 mortality during the first year of the pandemic in the reference lockdown scenario, as a fraction of COVID-19 mortality in the no-intervention scenario. Each dot represents a country and the color indicates the income group of the country: Low Income (blue), Lower-Middle-Income (green), Upper-Middle-Income (pink), and High-Income (red). Source: World Development Indicators, Penn World Tables, and the authors' calculations.

FIGURE 1

IMPACT OF THE REFERENCE LOCKDOWN ON ADULT COVID-19 MORTALITY



(a) Number of child deaths per COVID-19 fatality averted by the reference lockdown

(b) Total reduction in mortality caused by the reference lockdown

NOTES: Panel (a) presents the expected number of children lives lost per COVID-19 fatality averted against the logarithm of PPP-adjusted per capita GDP. Both the expected number of lives lost and the averted COVID-19 fatality are the differences between the reference lockdown policy and the no-intervention policy during the first year of the pandemic. Panel (b) presents the total reduction in mortality in the reference lockdown scenario, as a fraction of mortality in the no-intervention scenario. Each dot represents a country and the color indicates the income group of the country: Low Income (blue), Lower-Middle-Income (green), Upper-Middle-Income (pink), and High-Income (red). Source: World Development Indicators, Penn World Tables, and the authors' calculations.

FIGURE 2

IMPACT OF THE REFERENCE LOCKDOWN ON TOTAL MORTALITY

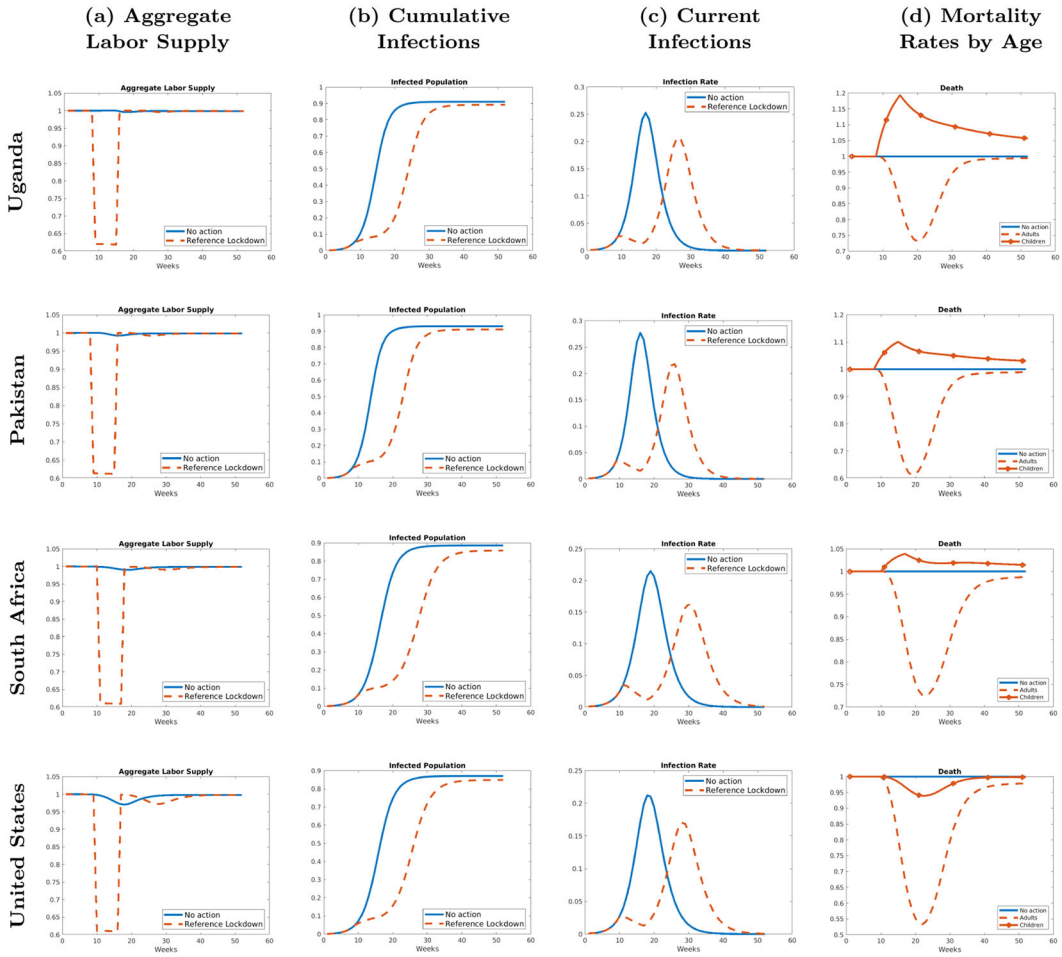
leads to more children's lives lost than COVID-19 fatalities averted. In the low-income country group, the reference lockdown causes an average of 1.76 child deaths per COVID-19 fatality averted. This rate is 0.59 in lower-middle-income countries, and 0.06 in upper-middle-income countries.

4.1.2. *Lockdowns and total mortality.* Another informative statistic to compute is the effect of the reference lockdown on total mortality, implicitly putting equal weight on every life lost, irrespective of age. Panel (b) of Figure 2 plots the reduction in total excess mortality achieved by the lockdown, relative to excess mortality in the no-action scenario against log per capita income. The highest average reduction in mortality (6.2%) is achieved in high-income countries where the lockdown prevents the most COVID-19 deaths and does not impact child mortality. For low- and middle-income countries, the net reductions in total mortality are smaller in magnitude as the lockdown both has less impact on COVID-19 mortality, and induces an increase in child mortality. In upper-middle- and lower-middle-income countries, mortality is reduced on average by 5.2% and 2%, respectively. In low-income countries, excess mortality increases by 2.6% with the lockdown since the economic contraction leads to a higher number of child deaths than the number of adult fatalities averted by the lockdown.

4.2. *Understanding the Intergenerational Mortality Trade-Offs.* The previous section illustrated the large variation in outcomes across countries following the reference lockdown. This section investigates the contributions of various country characteristics to the spread of the infection and subsequent mortality, both COVID-19-related and not.

4.2.1. *Lockdown and the dynamics of the COVID-19 pandemic.* To illustrate further what drives cross-country differences in outcomes, we present a more detailed analysis from four countries at different stages of economic development. We purposefully selected one country from each income group: Uganda (low income), Pakistan (lower-middle income), South Africa (upper-middle income), and the United States (high income). These different income levels dictate how consumption shortfalls due to lockdown policies would affect child mortality. The selected countries also differ substantially along other dimensions that determine the effect of lockdown policies, such as the population age distribution and health system capacity. Forty-six percent of the Ugandan population is under the age of 15, whereas only 3% are 60 years or older. In Pakistan (South Africa), 34 (29)% of the population are under 15 and 7 (9)% are 60 or older. The United States has the oldest population out of the four countries, with only 18% under 15% and 23% over 60. Uganda and Pakistan have only 0.5 and 0.6 hospital beds per 1,000 people to contrast with rates in South Africa and the United States of 2.8 and 2.9, respectively.

Column (a) of Figure 3 displays the aggregate labor supply during the first year of the pandemic as a fraction of the no-pandemic steady-state labor supply. The blue line represents labor supply without any government intervention and the red line represents the reference lockdown. Without a lockdown, there would be only small declines in labor supply during the weeks with the highest current infection rates (depicted in Panel (c) of the figure, solid blue line). This drop is entirely due to households limiting their own labor supply to lower COVID-19 transmission risks to their own members. Relative to the other countries, the drop in labor supply is largest in the United States given its substantially larger share of older adults in the population. However, even in the United States, the max of the labor decline in the no-action scenario is less than 5%. This muted response reflects the sizeable externality associated with pandemics, that is, households consider the trade-off between their members' mortality risk and income loss but not the impact their exposure could have on the further spread of the virus in the population. Under the lockdown scenario, there will be a uniform reduction of 38% in labor supply during the weeks in which the labor tax will be in effect. Then there are subsequent additional small reductions in labor supply when active infections



NOTES: This figure presents several pandemic-related indicators for selected countries under the no action scenario (solid blue line) and the reference lockdown scenario (dashed red line). Column (a) presents the change in aggregate labor supply, relative to the pre-COVID-19 steady state. Column (b) presents the cumulative infection rate, where the total population is normalized to 1. Column (c) shows the contemporaneous infection rate in each week. Column (d) portrays the cumulative all-cause mortality rates (from both COVID-19 and non-COVID-19) separately for children and adults under the reference lockdown relative to the no intervention policy. Source: authors' calculations.

FIGURE 3

SELECTED PANDEMIC INDICATORS: NO ACTION AND REFERENCE LOCKDOWN

reach their highest rates. As in the no-action scenario, the subsequent reduction in labor supply is largest in the United States but never exceeds 5%.

Columns (b) and (c) of Figure 3 illustrate how the lockdown policy affects virus transmission in the different countries. As can be seen in column (b), the reference lockdown will have only a negligible impact on the share of the population ever-infected by the end of the pandemic's first year. Instead, the lockdown slows the pace of transmission and displaces the peak infectivity period to later in the year. Overall, this policy slows the spread of the virus more effectively in wealthier countries. The infection rate peaks in Uganda and Pakistan before it does in South Africa and the United States. The primary reason for that is that the share of working-age adults in the total population is larger in the wealthier countries, and hence, the reduced economic activity has a larger impact on transmission rates in South Africa and the United States than it does in the other two countries.

Column (d) of Figure 3 depicts the cumulative all-cause child and adult mortality in the reference lockdown scenario, relative to the cumulative mortality in the no-lockdown scenario. In the three low- and middle-income countries, lockdown policies increase child mortality. This increase in child mortality is entirely due to the impact of the economic contraction induced by the lockdown. Given the high sensitivity of survival rates to income fluctuations in low-income countries, the largest increase in child mortality is in Uganda. In the United States, however, the lockdown policy reduces child mortality. Here, this reduction is entirely due to reduced COVID-19 child mortality, albeit from a low reference level.

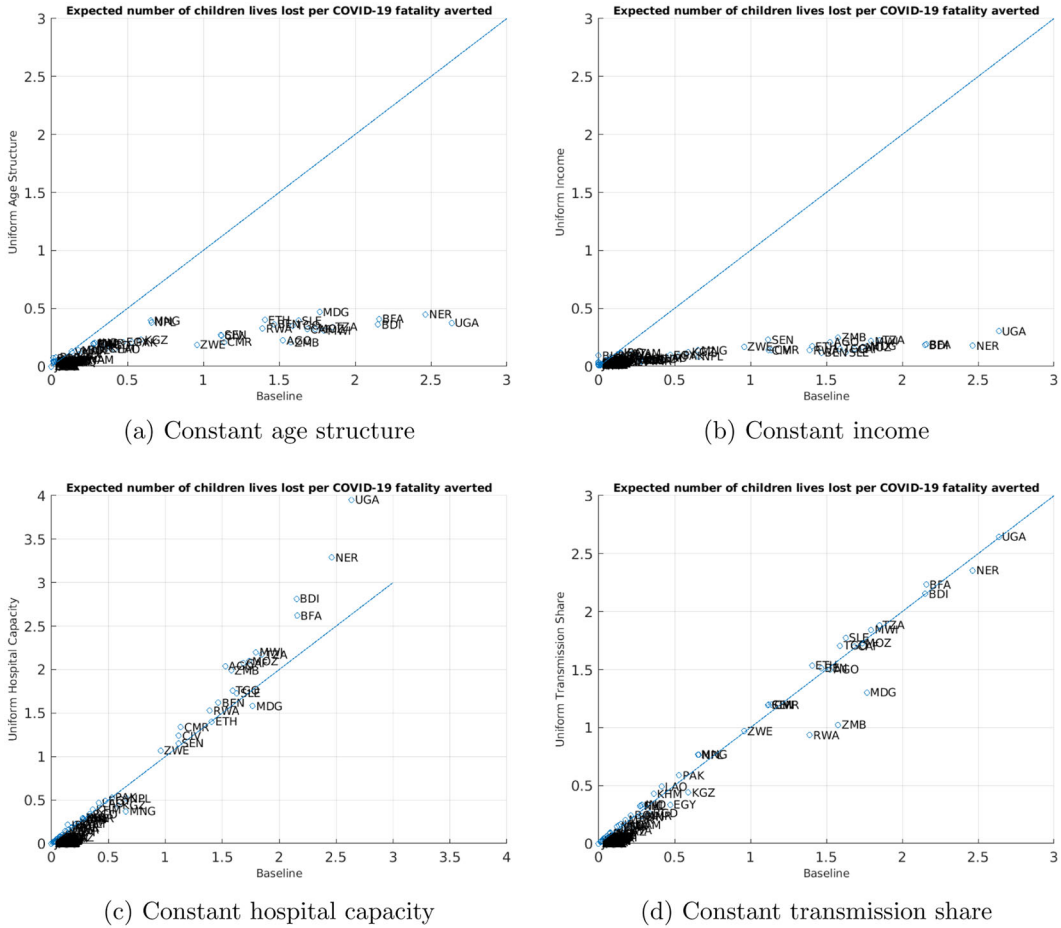
With respect to adults, the reference lockdown temporarily reduces mortality in all countries. However, by the end of the first year of the epidemic, the lockdown will have a small effect on the cumulative adult mortality, as already shown in Figure 1. As highlighted above, the lockdown slows the spread of the virus by a number of weeks but has only a minor impact on the cumulative rate of infections at the one-year horizon. Of the four countries, the 2% adult mortality reduction experienced in the United States is the biggest due to several factors. First, the United States has the highest share of adults over 60 who are at greater risk of COVID mortality. Second, the lockdown is most effective in slowing down the virus spread in the United States because of the differential modes of transmission. Third, because of higher hospital capacity, the slowdown in the virus spread causes a bigger improvement in survival rates.

Linking back to the results presented in Figure 2, in the poorest country of the four, Uganda, the total mortality in the lockdown scenario is higher than the no-action mortality rate by the end of the year. That is, the number of children who die from the GDP decline is greater than the number of COVID-19 deaths averted by the lockdown. In Pakistan, the excess child mortality is just slightly smaller than the modest adult mortality reduction achieved by the lockdown. In South Africa and the United States, the lockdown achieved positive although small reductions in total deaths.

4.2.2. Decomposing the heterogeneity in policy impact. This subsection presents counterfactual simulations designed to isolate various contributing mechanisms and gauge their influence on the overall cross-country variation in lockdown impacts.

4.2.3. Population age distribution. In the first simulation, we impose the same age distribution on all countries, equal to the unweighted average age distribution among the 85 sample countries. Figure 4(a) plots the number of children's lives lost per COVID-19 fatality averted by the reference lockdown in this counterfactual (y-axis) against the baseline. Relative to the results presented in Figure 2(a), the ratios are substantially lower when equalizing the age distribution across countries. The counterfactual death ratio is below the 45-degree line and below 0.5 for all countries. This suggests that the variation in the age distribution plays a crucial role in determining how the lockdowns affect overall mortality and the number of child deaths per COVID-19 fatality averted. Had the poorer countries had fewer children per adult, the number of child deaths per COVID-19 fatality averted would have been much lower country income.

In the second exercise, we keep the age structure of each country as in the data but assign all countries the same per capita income and therefore the same semielasticity of child mortality with respect to income. The income level in each country in this example is the geometric average of per capita incomes in the sample, corresponding to a level within the upper-middle-income designation. Figure 4(b) shows that the variation in children's lives lost per COVID-19 fatality averted shrinks considerably in this counterfactual. This implies that the cross-country differences in income are even more important than age structure in determining the impact of lockdown policies on overall mortality, given the relationship between income shortfalls and child mortality in poorer countries.



NOTES: This figure presents the expected number of children’s lives lost per COVID-19 fatality averted in the baseline scenario (x -axis) against four counterfactual scenarios on the y -axis. In Panel (a), the counterfactual imposes an identical age structure on all countries. In Panel (b), the counterfactual imposes the same income (belonging in the upper middle-income range) on all countries, and thus the same semielasticity of child mortality with respect to GDP fluctuations. In Panel (c), the counterfactual imposes the hospital capacity of the United States, measured as the number of hospital beds per thousand people, on all countries. In Panel (d), the counterfactual imposes the same π_{I1} , π_{I2} , and π_{I3} parameters on all countries, calibrated to the same targets as the United States. The solid blue line is the 45-degree line.


Source: Authors’ calculations.

FIGURE 4

COUNTERFACTUAL DEMOGRAPHICS, INCOME, HOSPITAL CAPACITY, AND TRANSMISSION SHARES

4.2.4. *Hospital capacity.* In the next counterfactual, we impose the U.S. hospital capacity on all countries. Figure 4(c) shows that the ratio of children’s lives lost per COVID-19 fatality averted increases for most low-income countries in this counterfactual. The reason is that with larger health systems, there are fewer COVID-19 fatalities under both the no-intervention and lockdown policies, leading to a smaller number of COVID-19 fatalities averted by the lockdown. It should be noted that in this simulation, the health system capacity only affects COVID-19-related mortality. Improved health system capacity may also reduce non-COVID-related mortality and improve child survival resilience to income shocks, but these channels are not incorporated in our model.

4.2.5. *COVID-19 transmission shares by activity.* Finally, Figure 4(d) depicts a counterfactual in which transmission probabilities in all the countries are calibrated such that the

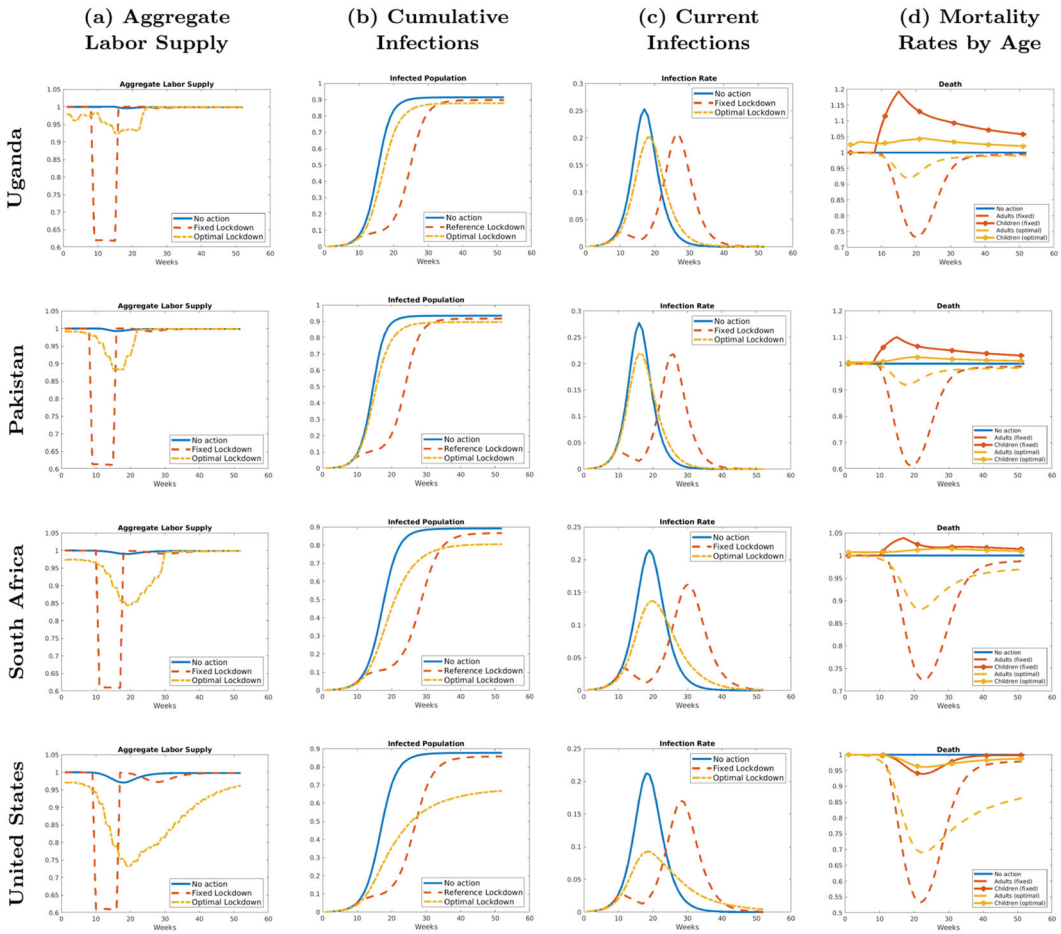


NOTES: This figure presents the expected number of children’s lives lost per COVID-19 fatality averted against the logarithm of PPP-adjusted per capita GDP. Both the expected number of lives lost and the averted COVID-19 fatalities are the differences between the optimal lockdown policy and the no-intervention policy during the first year from the beginning of the pandemic. Each dot represents a country and the color indicates the income group of the country: Low Income (blue), Lower-Middle-Income (green), Upper-Middle-Income (pink), and High-Income (red). Source: World Development Indicators, Penn World Tables, and the authors’ calculations.

FIGURE 5
NUMBER OF CHILD DEATHS PER COVID-19 FATALITY AVERTED BY THE OPTIMAL LOCKDOWN

share of transmission through each channel is similar to the United States. As explained above, the spread of COVID-19 in high-income countries is more reliant on work- and consumption-related activities, whereas it is more dependent on community transmission in developing countries. Therefore, the reference lockdown is more effective in slowing transmission in high-income countries, even though the reduction in aggregate labor supply induced by the policy is identical in all countries. Attributing the U.S. transmission parameters to all countries leads to only small declines in the ratio of children’s lives lost per COVID-19 fatality averted in poorer countries. This is because the change in transmission probabilities affects the distribution of infections among the different age groups regardless of whether a lockdown is imposed. Because only working-age adults both supply labor and conduct consumption-related activities in our model, increasing the weight of these channels implies that a larger share of the initial infections would be among this group. Therefore, at the time the reference lockdown is imposed, a smaller share of the elderly are infected, and fewer deaths would be averted by the lockdown. This increases the mortality ratio and explains why changing the transmission probabilities could increase the mortality ratio for some countries or only produce a moderate reduction in others.¹²

¹² A complementary approach to identifying the main mechanisms that drive cross-country variation in the number of child lives lost per COVID-19 fatality averted is a Oaxaca–Blinder decomposition that includes all the factors discussed in this subsection. This exercise is described in Appendix A.4. The two factors that explain most of the cross-country variation are: (a) differences across income groups in the semielasticity of child mortality to economic contraction, and (b) the share of national population below 15 years of age. A third significant factor is the importance of community-based transmission—when community transmission is relatively important, any lockdown policy will be less effective in averting COVID-19 fatalities.



NOTES: This figure presents several pandemic-related indicators for selected countries under the no action scenario (solid blue line), the reference lockdown scenario (dashed red line), and the optimal lockdown (dotted yellow line). Column (a) presents the change in aggregate labor supply, relative to the pre-COVID-19 steady state. Column (b) presents the cumulative infection rate, where the total population is normalized to 1. Column (c) shows the contemporaneous infection rate in each week. Column (d) portrays the cumulative all-cause mortality rates (from both COVID-19 and non-COVID-19) separately for children and adults under the reference lockdown relative to the no intervention policy.

Source: Authors' calculations.

FIGURE 6

SELECTED PANDEMIC INDICATORS: OPTIMAL LOCKDOWN

4.3. *An Optimal Lockdown Policy.* To conclude our discussion, we consider alternative lockdown policies that explicitly weigh COVID-19-related mortality against welfare more generally.

4.3.1. *Definitions and mortality differences.* We define an optimal lockdown policy as a labor tax sequence $\{\mu_t\}$ that maximizes the present value of aggregate social welfare, that is,

$$\max_{\{\mu_t\}} \int_{j=0}^1 U_j(\{\mu_t\})dj.$$

As such, the objective function now captures the trade-off between COVID-19 deaths and both increased infant mortality and the welfare loss due to reduced consumption. As this

problem does not yield a straightforward optimality condition, we use global maximization methods to search for the optimal lockdown policy.

Figure 5 depicts the ratio of child deaths per adult death averted by the optimal lockdown. Compared to the rates under the reference lockdown (Figure 2a), the ratio of child to adult mortality under the optimal policy is substantially lower. The ratio for all countries is below 0.7 and Uganda is the only country with a ratio above 0.5. Thus, in contrast to the reference lockdown calibrated to mimic policies implemented by European governments in the first few months of the pandemic, the optimal lockdown never leads to a net mortality increase.

Figure 6 demonstrates how the optimal lockdown policies vary across the four selected countries. Relative to the reference lockdown, the labor contraction in the optimal lockdown is smaller but starts earlier and lasts longer in all four countries. Column (a) of the figure also highlights the substantial differences in the length and severity of the optimal lockdowns among countries. There is a negative relation between a country's income level and the drop in labor supply under the optimal policy. Relative to the poorer countries, the lockdown in the United States is more severe as it has no impact on child mortality and is more effective in reducing transmission. In the United States, lockdown measures will be applied during the whole duration of the first year of the pandemic and the labor supply will decline by more than 25% when the current infection rate peaks. At the other extreme, Uganda would introduce lockdown measures only in the first half of the year and labor supply never drops below 90% of the prepandemic steady state.

As seen in column (b) of Figure 6, the optimal lockdown policy substantially reduces the share of the population that ever gets infected in South Africa and the United States but not in Uganda and Pakistan. As a result, the optimal policy has a much larger impact on adult mortality in the wealthier countries, in comparison to the reference lockdown (column (d)). On the other hand, the optimal policy induces smaller increases in child mortality in the low- and middle-income countries.

5. CONCLUSION

At the start of the COVID-19 pandemic, countries around the world imposed lockdown measures similar in severity. Our analysis, however, suggests that optimal policies substantially differ, depending on the vulnerability of child survival to income shocks, countries' demographic characteristics, and patterns of social contacts. The reason is that economic contractions in low- and middle-income countries increase child mortality, and policy responses to the pandemic contributed to declines in national income in much of the world. This article highlights and then quantifies this relatively neglected consequence of lockdowns purely in terms of intergenerational mortality trade-offs, thereby informing country-specific assessments of the costs and benefits of lockdowns as policies to fight a pandemic.

Since our main objective is the formulation of a new trade-off, we abstracted from other channels through which lockdowns might affect health outcomes. Importantly, the impacts of such channels might differ across countries, thereby generating heterogeneity in how they affect the choice of lockdown policies. The simulations in our article should therefore not be a source of definitive policy prescriptions, but rather a quantification of the importance of an intergenerational trade-off that had been largely overlooked in both academic and policy fora.

An important lesson from this analysis, which holds beyond the COVID-19 pandemic, is that nonpharmaceutical interventions to prevent infectious disease spread will likely affect subgroups of the population in an adverse manner and will do so in a context-specific way. This implies that additional targeted policy instruments could be used to alleviate the downside distributional effects. In our setting, targeted income transfers to households with young children and pregnant mothers are examples of mitigating policies that can be adopted during lockdowns. Given the positive externality that epidemic containment policies have on the rest of the world (in part because of the lower likelihood of the emergence of novel and possibly

more lethal variants), there is a rationale for these aforementioned mitigating instruments to be financed through development assistance to low-income countries.

APPENDIX A

A.1 Model and Quantification.

A.1.1 Proof of Lemma 1. First, note that we can divide the future states starting from any k into two categories: those k' that depend on n_{kt} and those that do not. We define the set of future states to which the transition probability that are independent from the household decision as \bar{K}_k :

$$\bar{K}_k = \{k' | \rho_{kk'}(n_{kt} | \Theta_t) = \rho_{kk'}(\Theta_t)\}.$$

All the infeasible states from k with $\rho_{kk'} = 0$ are also part of the set \bar{K}_k . With this insight, we can rewrite the household decision problem, highlighting that from the household's perspective, the transition probabilities into the states $k' \in \bar{K}_k$ will be taken as exogenously given:

$$(A.1) \quad U_k = u(n_{kt}) + \beta \left[\sum_{k' \in (K \setminus \bar{K}_k)} \rho_{kk'}(n_{kt} | \Theta_t) U_{k'} + \sum_{k' \in \bar{K}_k} \rho_{kk'}(\Theta_t) U_{k'} \right].$$

The household will internalize the dynamic effects of labor supply on (i) child mortality and (ii) virus transmission through Equations (15) and (16). In the rest of the section, we discuss all the possible combinations of these two dynamic effects.

Case 1, Infection Shock ($\pi_{Ikt}(2)$) Applicable. In the first case, the household has a susceptible young adult but no children, and thus, only the infection shock applies. We express the household problem as:

$$(A.2) \quad \begin{aligned} U_k &= u(n_{kt}) + \beta \pi_{Ikt}(2)(1 - \pi_{nkt}(2)) \sum_{k'(2)=I} \rho_{kk'}(1 | \Theta_t) \rho_{kk'}(3 | \Theta_t) U_{k'} \\ &+ \beta(1 - \pi_{Ikt}(2))(1 - \pi_{nkt}(2)) \sum_{k'(2)=S} \rho_{kk'}(1 | \Theta_t) \rho_{kk'}(3 | \Theta_t) U_{k'} \\ &+ \sum_{k' \in \bar{K}_k} \rho_{kk'}(\Theta_t) U_{k'}. \end{aligned}$$

In the equation above, we separate the future states depending on the outcome of the young adult, $k'(2)$, which relies on the realization of $\pi_{Ikt}(2)$. We also use $\rho_{kk'}(m)$ to denote the probability that the state of individual m will change from $k(m)$ to $k'(m)$ in the next period.¹³

$$\rho_{kk'}(n_{kt} | \Theta_t) = \prod_{m=1}^3 \rho_{kk'}(m | \Theta_t).$$

¹³ The equation above implicitly relies on the assumption that the transition shocks are independent across family members. Note that the future states in which the young adult dies due to non-COVID reasons, $k'(2) = D$, are part of the set \bar{K}_k , and thus, do not affect the policy function of the household.

Considering the transition probabilities, the Lagrangian of the household problem is:

$$\begin{aligned}
 U_k &= u(n_{kt}) + \beta\pi_{Ikt}(2)(1 - \pi_{nkt}(2)) \sum_{k'(2)=I} \rho_{kk'}(1|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} \\
 &+ \beta(1 - \pi_{Ikt}(2))(1 - \pi_{nkt}(2)) \sum_{k'(2)=S} \rho_{kk'}(1|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} + \sum_{k' \in (\bar{K}_k)} \rho_{kk'}(\Theta_t)U_{k'} \\
 &- \lambda_{\pi_t}^k [\pi_{I1}[(1 - \mu_t)w_t n_{kt} + \Gamma_{kt}]C_{It} + \pi_{I2}n_{kt}N_{It} + \pi_{I3}I_t + \pi_{I4}\mathbf{I}_k - \pi_{Ikt}(2)],
 \end{aligned}$$

where $\lambda_{\pi_t}^k$ is the multiplier associated with the transition probability. The first-order condition (FOC) with respect to n_{kt} is

$$(A.3) \quad \frac{\partial u}{\partial n_{kt}} - \lambda_{\pi_t}^k [\pi_{I1}(1 - \mu_t)w_t C_{It} + \pi_{I2}N_{It}] = 0.$$

Finally, the FOC for $\pi_{Ikt}(2)$ is:

$$(A.4) \quad \beta(1 - \pi_{nkt}(2)) \left\{ \sum_{k'(2)=S} \rho_{kk'}(1|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} - \sum_{k'(2)=I} \rho_{kk'}(1|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} \right\} = \lambda_{\pi_t}^k.$$

Note that from the equation above, it is clear that $\lambda_{\pi_t}^k$ captures the increase in the future value if the young adult avoids infection in the next period. In other words, $\lambda_{\pi_t}^k$ is the value of staying uninfected during the pandemic.

Case 2, Child Mortality Shock ($\pi_{dkt}(1)$) Applicable. In this case, the young working adult is either infected or recovered, and the children are alive. The Lagrangian equation is:¹⁴

$$\begin{aligned}
 U_k &= u(n_{kt}) + \beta\pi_{dkt}(1) \sum_{k'(1)=D} \rho_{kk'}(2|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} \\
 &+ \beta(1 - \pi_{dkt}(1)) \sum_{k'(1) \neq D} \rho_{kk'}(2|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} + \sum_{k' \in (\bar{K}_k)} \rho_{kk'}(\Theta_t)U_{k'} \\
 &- \lambda_{\pi_d}^k \left\{ \bar{\pi}_n(m) + v(m) \left[1 - \frac{(1 - \mu_t)w_t n_{kt} + \Gamma_{kt}}{\tilde{c}_k} \right] - \pi_{dkt}(1) \right\},
 \end{aligned}$$

The FOC with respect to n_{kt} is:

$$(A.5) \quad \frac{\partial u}{\partial n_{kt}} + \lambda_{\pi_d}^k v(1) \frac{(1 - \mu_t)w_t}{\tilde{c}_k} = 0,$$

and the FOC with respect to $\pi_{dkt}(1)$:

$$(A.6) \quad \beta \left\{ \sum_{k'(1) \neq D} \rho_{kk'}(2|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} - \sum_{k'(1)=D} \rho_{kk'}(2|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} \right\} = \lambda_{\pi_d}^k.$$

¹⁴ Note that in the equation above, we have assumed that the children are not infected. Assuming infected children will only alter the level of $\pi_{dkt}(1)$ without affecting the first order conditions. The FOC also assumes that the young working adult is not infected. If the adult is infected, the FOC will be slightly modified to reflect the discount on productivity.

Similar to the previous case, the multiplier $\lambda_{\pi_d}^k$ captures the value of keeping the children alive.

Case3, Both Shocks Applicable. When both the infection and the mortality shocks apply, the Lagrangian becomes:

$$\begin{aligned}
 U_k = & u(n_{kt}) + \beta\pi_{dkt}(1)\pi_{lkt}(2)[1 - \pi_{nkt}(2)] \sum_{k'(1)=D \cap k'(2)=I} \rho_{kk'}(3|\Theta_t)U_{k'} \\
 & + \beta[1 - \pi_{dkt}(1)]\pi_{lkt}(2)[1 - \pi_{nkt}(2)] \sum_{k'(1) \neq D \cap k'(2)=I} \rho_{kk'}(3|\Theta_t)U_{k'} \\
 & + \beta\pi_{dkt}(1)[1 - \pi_{lkt}(2)][1 - \pi_{nkt}(2)] \sum_{k'(1)=D \cap k'(2)=S} \rho_{kk'}(3|\Theta_t)U_{k'} \\
 & + \beta[1 - \pi_{dkt}(1)][1 - \pi_{lkt}(2)][1 - \pi_{nkt}(2)] \sum_{k'(1) \neq D \cap k'(2)=S} \rho_{kk'}(3|\Theta_t)U_{k'} \\
 & + \beta\pi_{dkt}(1)\pi_{nkt}(2) \sum_{k'(1)=D \cap k'(2)=D} \rho_{kk'}(3|\Theta_t)U_{k'} \\
 & + \beta[1 - \pi_{dkt}(1)]\pi_{nkt}(2) \sum_{k'(1) \neq D \cap k'(2)=D} \rho_{kk'}(3|\Theta_t)U_{k'} + \sum_{k' \in (\bar{K}_k)} \rho_{kk'}(\Theta_t)U_{k'} \\
 & - \lambda_{\pi_l}^k [\pi_{l1}[(1 - \mu_t)w_t n_{kt} + \Gamma_{kt}]C_{lt} + \pi_{l2}n_{kt}N_{lt} + \pi_{l3}I_t + \pi_{l4}\mathbf{I}_k - \pi_{lkt}(2)] \\
 & - \lambda_{\pi_d}^k \left\{ \bar{\pi}_n(m) + v(m) \left[1 - \frac{(1 - \mu_t)w_t n_{kt} + \Gamma_{kt}}{\tilde{c}_k} \right] - \pi_{dkt}(1) \right\}.
 \end{aligned}$$

The FOC for labor supply is:

$$(A.7) \quad \frac{\partial u}{\partial n_{kt}} - \lambda_{\pi_l}^k [\pi_{l1}(1 - \mu_t)w_t C_{lt} + \pi_{l2}N_{lt}] + \lambda_{\pi_d}^k v(1) \frac{(1 - \mu_t)w_t}{\tilde{c}_k} = 0$$

and for $\pi_{lkt}(2)$:

$$(A.8) \quad \beta(1 - \pi_{nkt}(2)) \left\{ \sum_{k'(2)=S} \rho_{kk'}(1|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} - \sum_{k'(2)=I} \rho_{kk'}(1|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} \right\} - \lambda_{\pi_l}^k = 0.$$

Finally, the FOC for $\pi_{dkt}(1)$:

$$(A.9) \quad \beta \left\{ \sum_{k'(1) \neq D} \rho_{kk'}(2|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} - \sum_{k'(1)=D} \rho_{kk'}(2|\Theta_t)\rho_{kk'}(3|\Theta_t)U_{k'} \right\} - \lambda_{\pi_d}^k = 0.$$

Case 4, Neither Shocks Applicable. In this case, the labor decision today does not affect the transition probability. As a result, the FOC of labor supply only takes the current utility into consideration:

$$\frac{\partial u}{\partial n_{kt}} = 0.$$

Case 5, No Young Adult. In the four cases above, we have implicitly assumed that the young adult is alive. In the case of a deceased young adult, there is no optimization of

problem for the household to solve. The remaining members of the household will consume the government transfer in every period.

A.1.2 Recursion. We use recursion to compute the postpandemic steady-state levels of U_k , as well as U_{kt} in the backward induction.

To infer the steady state U_k , the necessary condition is that $U_{k'}$ are known for all the $k' \neq k$ such that $\rho_{kk'} > 0$. In other words, we need to know the value of all the future states k' that state k can possibly transit into, except for k itself, to infer the value of state k . We start the recursion by assuming that no U_k is known. In the first iteration, the only state that can be inferred is the absorbing state, $k = 64$, in which all the agents are dead. In the second recursion, we can then infer all the states that are only one step away from the absorbing state, $k = 61, 62, 63$. We repeat this process until all the U_k are known. The last state to infer is $k = 1$.

With the steady state U_k computed for all k , we then repeat this process for all $t = \mathcal{T}, \mathcal{T} - 1, \dots, 1$ to compute all the U_{kt} . At $t = \mathcal{T}$, we assume that the future states are steady state. For all $t < \mathcal{T}$, the future states are simply $t + 1$.

A.1.3 Calibration of σ . From Equation (20), it is straightforward to see that in the prepandemic steady state, $\Lambda_{kt} = 0$ and $\mu_t = 0$. Therefore, the labor supply is reduced to:

$$(A.10) \quad n_k = \frac{\left[\sum_{m=1}^3 \ell^m \mathbf{1}_k(m) \right]^{\frac{1}{\sigma-1}} w_t}{\theta}.$$

The instantaneous utility function in Equation (17) simplifies to:

$$\begin{aligned} u(n_k) &= (w_t n_k) \left[\sum_{m=1}^3 \ell^m \mathbf{1}_k(m) \right]^{\frac{1}{\sigma-1}} - \frac{\theta}{2} n_k^2 \\ &= \left(\frac{\left[\sum_{m=1}^3 \ell^m \mathbf{1}_k(m) \right]^{\frac{1}{\sigma-1}} w_t}{\theta} \right) \left[\sum_{m=1}^3 \ell^m \mathbf{1}_k(m) \right]^{\frac{1}{\sigma-1}} - \frac{\theta}{2} \left(\frac{\left[\sum_{m=1}^3 \ell^m \mathbf{1}_k(m) \right]^{\frac{1}{\sigma-1}} w_t}{\theta} \right)^2 \\ &= \frac{\left[\sum_{m=1}^3 \ell^m \mathbf{1}_k(m) \right]^{\frac{2}{\sigma-1}} w_t^2}{2\theta}. \end{aligned}$$

From the expression above, it is straightforward to see that when $\sigma = 3$, the death of a non-productive member with mass ℓ^m reduces $u(n_k)$ by ℓ^m fraction.

A.2 Robustness Checks.

A.2.1 Simple adjustments to the reference lockdown. Although the optimal lockdown involves a high-dimensional global optimization problem, an interesting question is whether governments could have done better than the reference lockdown by varying either the duration or the onset timing of the reference lockdown. Appendix Table A.6 investigates this. It summarizes our key statistic, the ratio of children’s deaths per adult COVID-19 fatality averted under alternative reference lockdown scenarios. The middle panel shows what happens when we lengthen the reference lockdown between two and eight weeks. If anything, for the lower-income countries, the mortality ratio rises under longer lockdowns, as longer lockdowns have a disproportionately larger effect on child mortality relative to adult lives saved.

On the other hand, delaying the lockdown onset lowers the mortality ratio substantially, putting it firmly below 1. Evidently, delaying the lockdown is beneficial. Looking closer, it

appears that the fall in the mortality ratio is driven entirely by the number of adult lives saved. The second and third columns of the table report the number of children's deaths and the number of adult lives saved in each alternative scenario, normalized by the number of deaths in the baseline reference lockdown. Delaying the lockdown by two to eight weeks hardly changes the number of child deaths. This is not surprising, as the child mortality comes from the economic downturn engineered by the lockdown, and so, it does not matter very much whether the economic downturn occurs now or eight weeks from now. However, the third column shows that delaying the lockdown leads to many more adult lives saved, compared to the baseline lockdown onset, due to a flattening of the epidemic peak that renders health system constraints less binding. In other words, a lockdown is more effective when the number of active infections is higher, and the hospital capacities are stretched. Overall, the results do not change qualitatively. There is a substantial number of children's lives lost per COVID-19 fatality averted in low-income countries in all hypothetical scenarios.

A.2.2 Valuing household members: σ . One parameter that is challenging to discipline is σ , the elasticity of substitution between consumptions of household members. It regulates how much the utility of the household drops when it loses a member, which, in turn, conditions the labor supply of the working adult. Our baseline approach picks σ so that the reduction in the household's utility from a family member's death equals the share of that family member in the population of the household (see Appendix A.1.3). This has a natural social planner interpretation, as it corresponds to a "utilitarian" objective function that puts the same Pareto weight on every member of society.

An alternative approach is to appeal to family economics by using the fertility literature to infer σ . To do this, we used a simplified version of the fertility choice model of Doepke (2004) to infer the value of children in household utility for each country based on per capita consumption and total fertility data from the World Development Indicators (WDI). In our sample of 85 countries, the average value of children is around 37% of household utility. We abstract from the distinction between skilled and unskilled children in Doepke (2004) and compute the value of children as a fraction of household utility as $0.132 * \sqrt{n} / (\sqrt{c} + 0.132 * \sqrt{n})$, where n is the number of children and c is the consumption. We approximate the number of children using the total fertility rate and consumption using the per capita consumption in WDI. The value of children as a fraction of household utility in our model is computed as

$$1 - \frac{u_{(k=4)}(\sigma)}{u_{(k=1)}(\sigma)},$$

where $u_{(k)}$ is the pre-pandemic steady-state level of utility of a household in state k ; $k = 1$ is the state in which all the members are alive, and $k = 4$ is the state in which the children are deceased (see Appendix Table A.1 for the list of household states). This exercise yields $\sigma = 2.5$.

Finally, we could assess sensitivity by treating σ as a free parameter. Appendix Table A.7 compares the basic results from using the baseline σ , the fertility literature-based σ , as well as a wider range of σ 's from 1 to 5. As expected, a lower σ implies that both the reference and the optimal lockdowns lead to fewer children's lives lost per COVID-19 fatality averted (top panel). This is because a lower σ implies a stronger private labor supply response, making the government-imposed lockdown less needed (bottom panel). Qualitatively, the results are unchanged. Even with σ very close to 1, there is a substantial number of children's lives lost per COVID-19 fatality averted in the lower-income countries. We cannot lower σ below 1 because as often happens in similar CES settings, the model is not continuous at $\sigma = 1$. Instead it is a hyperbola. Values of $\sigma < 1$ have some perverse implications, as they imply that household members' consumptions are complements. In this environment losing a family member actually raises household utility conditional on the same level of consumption expenditure—the mirror image of the "love for variety" effect. The value of σ has a limited

effect on the results because it affects the household’s valuation of both children’s and elderly lives. When deciding on how much labor to supply, a low-income country household balances the increased elderly mortality from working with the increased children’s mortality from not working. Lowering σ makes the loss of both elderly and children more painful for the household, but these risks counteract each other in the labor supply decision.

A.3 Additional Tables and Figures’.

TABLE A.1
HOUSEHOLD STATES

<i>k</i>	States	Case	<i>k</i>	States	Case	<i>k</i>	States	Case
1	SSS	3	23	ISR	3	45	IDR	5
2	ISS	3	24	RSR	3	46	IRD	2
3	RSS	3	25	DSR	1	47	IRR	2
4	DSS	1	26	ISD	3	48	IDD	5
5	SIS	2	27	RSD	3	49	DIR	4
6	SRS	2	28	DSD	1	50	RID	2
7	SDS	5	29	IIS	2	51	RIR	2
8	SSI	3	30	RIS	2	52	DID	4
9	SSR	3	31	DIS	4	53	DRI	4
10	SSD	3	32	IRS	2	54	RDI	5
11	SII	2	33	RRS	2	55	RRI	2
12	SRI	2	34	DRS	4	56	DDI	5
13	SDI	5	35	IDS	5	57	RRR	2
14	SIR	2	36	RDS	5	58	DRR	4
15	SRR	2	37	DDS	5	59	RDR	5
16	SDR	5	38	III	2	60	RRD	2
17	SID	2	39	RII	2	61	RDD	5
18	SRD	2	40	DII	4	62	DRD	4
19	SDD	5	41	IRI	2	63	DDR	5
20	ISI	3	42	IDI	5	64	DDD	5
21	RSI	3	43	IIR	2			
22	DSI	1	44	IID	2			

NOTE: This table lists all the states that a household could be in. The three letters indicate the state of each of the three members of the household. “S” means that the member is susceptible, “I” for infected, “R” for recovered, and “D” for deceased. For example, for a household in state 53, “DRI,” children are deceased, adults recovered, and the elderly infected. “Case” refers to the cases used to prove Lemma 1, as detailed in Appendix A.1.1.

TABLE A.2
MORTALITY AND HOSPITALIZATION RATES BY AGE

Age Group	Hospitalization	Mortality
0–14	0.0009	0.00003
15–60	0.023	0.001
> 60	0.130	0.034

NOTE: Average hospitalization and mortality rate by age group in low- and middle-income countries. Source: Walker et al. (2020).

TABLE A.3
LIST OF COUNTRIES

Low-income countries (15):			
Benin	Burkina Faso	Burundi	Central African Republic
Ethiopia	Madagascar	Malawi	Mozambique
Nepal	Niger	Rwanda	Sierra Leone
Tanzania	Togo	Uganda	
Lower-middle income countries (22):			
Angola	Bangladesh	Bolivia	Cambodia
Cameroon	Cote d'Ivoire	Egypt, Arab Republic	El Salvador
India	Indonesia	Kyrgyz Republic	Lao PDR
Mongolia	Morocco	Myanmar	Nicaragua
Pakistan	Philippines	Senegal	Vietnam
Zambia	Zimbabwe		
Upper-middle income countries (31):			
Albania	Algeria	Argentina	Armenia
Azerbaijan	Belarus	Belize	Bosnia and Herzegovina
Botswana	Brazil	Bulgaria	China
Colombia	Costa Rica	Dominican Republic	Ecuador
Fiji	Georgia	Iraq	Jamaica
Jordan	Kazakhstan	Lebanon	Malaysia
Mexico	Namibia	Paraguay	Peru
South Africa	Sri Lanka	Thailand	
High-income countries (17):			
Austria	Bahamas, The	Bahrain	Barbados
Belgium	Chile	Denmark	France
Germany	Italy	Japan	Norway
Panama	Spain	Switzerland	United Kingdom
United States			

NOTE: This table lists the 85 countries included in the analysis by income group classification based on the World Bank grouping for fiscal year 2020.

TABLE A.4
COUNTRY-LEVEL TABLE, PART 1

Country	$\pi_{I1} \times (\frac{A_i}{A_{USA}})^2$	π_{I2}	π_{I3}	ℓ^1	ℓ^2	ℓ^3	Beds	Start Week
AGO	0.278	0.306	0.415	0.464	0.499	0.037	0.8	9
ALB	0.182	0.207	0.343	0.172	0.616	0.212	2.9	12
ARG	0.191	0.216	0.348	0.244	0.600	0.155	5.0	12
ARM	0.195	0.221	0.366	0.208	0.607	0.185	4.2	11
AUT	0.317	0.285	0.261	0.144	0.599	0.257	7.6	12
AZE	0.192	0.216	0.393	0.235	0.648	0.116	4.7	11
BDI	0.274	0.302	0.418	0.453	0.507	0.041	0.8	9
BEL	0.389	0.342	0.306	0.170	0.574	0.256	6.2	10
BEN	0.264	0.292	0.431	0.419	0.530	0.051	0.5	9
BFA	0.270	0.298	0.426	0.444	0.517	0.039	0.4	9
BGD	0.215	0.241	0.461	0.268	0.653	0.080	0.8	9
BGR	0.190	0.216	0.324	0.147	0.571	0.282	6.8	13
BHR	0.282	0.255	0.340	0.183	0.764	0.053	2.0	10
BHS	0.297	0.269	0.285	0.216	0.662	0.122	2.9	12
BIH	0.183	0.208	0.333	0.145	0.602	0.253	3.5	13
BLR	0.189	0.214	0.345	0.172	0.601	0.226	11.0	12
BLZ	0.190	0.215	0.375	0.292	0.632	0.076	1.3	11
BOL	0.204	0.230	0.373	0.302	0.594	0.104	1.1	11

(Continued)

TABLE A.4
(CONTINUED)


Country	$\pi_{I1} \times (\frac{A_i}{A_{USA}})^2$	π_{I2}	π_{I3}	ℓ^1	ℓ^2	ℓ^3	Beds	Start Week
BRA	0.173	0.195	0.346	0.207	0.652	0.140	2.2	12
BRB	0.307	0.277	0.254	0.168	0.601	0.232	5.8	13
BWA	0.235	0.262	0.449	0.334	0.596	0.070	1.8	9
CAF	0.274	0.302	0.435	0.435	0.520	0.045	1.0	9
CHE	0.321	0.288	0.262	0.150	0.598	0.253	4.7	12
CHL	0.301	0.272	0.270	0.192	0.634	0.174	2.2	12
CHN	0.195	0.219	0.401	0.177	0.649	0.174	4.2	10
CIV	0.261	0.289	0.435	0.415	0.538	0.047	0.4	9
CMR	0.261	0.289	0.432	0.421	0.536	0.043	1.3	9
COL	0.177	0.200	0.352	0.222	0.647	0.132	1.5	12
CRI	0.174	0.197	0.341	0.208	0.641	0.150	1.2	12
DEU	0.372	0.329	0.291	0.140	0.574	0.286	8.3	11
DNK	0.332	0.298	0.257	0.163	0.576	0.261	2.5	12
DOM	0.193	0.218	0.365	0.274	0.615	0.111	1.6	11
DZA	0.204	0.230	0.371	0.308	0.593	0.099	1.9	11
ECU	0.192	0.217	0.365	0.274	0.616	0.110	1.5	11
EGY	0.239	0.267	0.439	0.339	0.579	0.082	1.6	9
ESP	0.323	0.290	0.261	0.144	0.593	0.263	3.0	12
ETH	0.255	0.283	0.435	0.399	0.548	0.053	0.3	9
FJI	0.200	0.225	0.382	0.290	0.614	0.096	2.3	11
FRA	0.436	0.377	0.328	0.177	0.555	0.268	6.5	9
GBR	0.367	0.325	0.291	0.177	0.579	0.244	2.8	11
GEO	0.199	0.225	0.352	0.202	0.583	0.215	2.6	12
IDN	0.197	0.221	0.399	0.259	0.640	0.101	1.2	10
IND	0.224	0.250	0.468	0.262	0.637	0.101	0.7	9
IRQ	0.202	0.228	0.347	0.377	0.572	0.051	1.4	12
ITA	0.331	0.296	0.252	0.130	0.572	0.298	3.4	12
JAM	0.181	0.205	0.352	0.234	0.634	0.133	1.7	12
JOR	0.199	0.224	0.376	0.329	0.611	0.061	1.4	11
JPN	0.340	0.303	0.228	0.124	0.532	0.343	13.4	13
KAZ	0.215	0.242	0.392	0.291	0.586	0.122	6.7	10
KGZ	0.216	0.243	0.400	0.326	0.593	0.081	4.5	10
KHM	0.195	0.220	0.372	0.309	0.615	0.076	0.8	11
LAO	0.196	0.222	0.374	0.319	0.613	0.068	1.5	11
LBN	0.196	0.220	0.393	0.251	0.637	0.112	2.9	11
LKA	0.231	0.258	0.443	0.237	0.599	0.164	3.6	9
MAR	0.202	0.227	0.388	0.268	0.614	0.119	1.1	11
MDG	0.254	0.282	0.434	0.401	0.549	0.050	0.2	9
MEX	0.187	0.211	0.363	0.258	0.629	0.112	1.5	11
MMR	0.196	0.220	0.401	0.255	0.645	0.100	0.9	10
MNG	0.197	0.223	0.378	0.311	0.616	0.073	7.0	11
MOZ	0.272	0.300	0.426	0.441	0.516	0.044	0.7	9
MWI	0.265	0.293	0.431	0.430	0.529	0.041	1.3	9
MYS	0.191	0.215	0.398	0.234	0.656	0.110	1.9	11
NAM	0.241	0.269	0.440	0.368	0.576	0.056	2.7	9
NER	0.303	0.330	0.405	0.497	0.462	0.041	0.3	9
NIC	0.195	0.220	0.373	0.295	0.618	0.087	0.9	11
NOR	0.332	0.297	0.272	0.173	0.595	0.233	3.9	12
NPL	0.229	0.256	0.469	0.288	0.625	0.087	0.3	9
PAK	0.248	0.275	0.465	0.348	0.584	0.067	0.6	9
PAN	0.338	0.302	0.292	0.265	0.613	0.122	2.3	11
PER	0.184	0.208	0.354	0.247	0.628	0.125	1.6	12
PHL	0.200	0.226	0.383	0.300	0.613	0.086	1.0	11
PRY	0.196	0.221	0.370	0.289	0.612	0.099	1.3	11
RWA	0.252	0.280	0.437	0.395	0.554	0.051	1.6	9
SEN	0.267	0.295	0.431	0.426	0.526	0.048	0.3	9

(Continued)

TABLE A.4
(CONTINUED)

Country	$\pi_{I1} \times (\frac{A_i}{A_{USA}})^2$	π_{I2}	π_{I3}	ℓ^1	ℓ^2	ℓ^3	Beds	Start Week
SLE	0.257	0.285	0.444	0.403	0.550	0.046	0.4	9
SLV	0.192	0.217	0.362	0.266	0.614	0.121	1.3	11
TGO	0.257	0.285	0.440	0.406	0.547	0.047	0.7	9
THA	0.195	0.220	0.397	0.166	0.642	0.192	2.1	11
TZA	0.267	0.295	0.426	0.436	0.522	0.042	0.7	9
UGA	0.272	0.300	0.416	0.460	0.508	0.032	0.5	9
USA	0.371	0.328	0.303	0.184	0.588	0.229	2.9	10
VNM	0.195	0.219	0.396	0.232	0.645	0.123	2.6	11
ZAF	0.198	0.223	0.389	0.288	0.627	0.085	2.8	11
ZMB	0.264	0.292	0.424	0.440	0.526	0.034	2.0	9
ZWE	0.266	0.294	0.439	0.419	0.535	0.046	1.7	9

NOTE: This table lists the following calibrated country-specific parameters: the infection rates ($\pi_{I(\cdot)}$), age structure ($\ell^{(\cdot)}$), hospital capacity measured in the number of beds per thousand population (“beds”), and the starting week of the reference lockdown (“start week”). Because the consumption-related infection probability $\pi_{I(\cdot)}$ is scaled by the product of consumptions $c_{jt}C_{It}$, and in each country productivity is normalized relative to the United States, to make the parameters comparable across countries we renormalize them by the square of the relative productivity of the country to the United States.



NOTES: This figure presents the expected number of COVID-19 cases averted by the reference lockdown against the logarithm of per capita GDP. The number of cases is normalized by the total population of each country. The expected number of cases averted is the difference between a lockdown policy and the no-intervention policy during the first year since the beginning of the pandemic. Each dot represents a country and the color indicates the income group of the country: Low Income (blue), Lower-Middle-Income (green), Upper-Middle-Income (pink), and High-Income (red).

Source: World Development Indicators, Penn World Tables, and the authors’ calculations.

FIGURE A.1

NUMBER OF COVID-19 CASES AVERTED BY THE REFERENCE LOCKDOWN, PER CAPITA

TABLE A.5
COUNTRY-LEVEL TABLE, PART 2

Country	$\pi_d(1) \times 10^5$	$\pi_d(2) \times 10^3$	$\pi_d(3) \times 10^2$	$\pi_b(1) \times 10^4$	$\pi_b(2) \times 10^2$	$\pi_b(3) \times 10^1$	$\bar{\pi}_n(1) \times 10^5$	$\bar{\pi}_n(2) \times 10^4$	$\bar{\pi}_n(3) \times 10^3$
AGO	3.180	1.092	3.409	8.392	1.806	1.263	13.218	1.164	1.845
ALB	3.220	1.891	3.919	8.496	2.658	1.321	2.001	0.428	1.450
ARG	3.220	1.574	4.256	8.505	2.385	1.357	1.674	0.506	1.128
ARM	3.200	1.857	3.835	8.460	2.687	1.291	2.059	0.548	1.415
AUT	3.220	2.142	4.398	8.501	2.974	1.371	0.577	0.256	0.927
AZE	3.150	1.795	3.499	8.325	2.631	1.256	4.023	0.536	1.532
BDI	3.170	1.037	3.288	8.365	1.751	1.243	14.166	1.471	1.976
BEL	3.260	2.121	4.364	8.615	2.929	1.363	0.539	0.283	0.918
BEN	3.190	1.161	3.608	8.418	1.887	1.290	16.085	1.183	1.880
BFA	3.190	1.089	3.366	8.426	1.808	1.262	14.767	1.268	2.534
BGD	3.240	1.444	3.950	8.556	2.249	1.324	4.929	0.593	1.196
BGR	3.270	2.096	4.194	8.644	2.992	1.355	1.424	0.614	1.470
BHR	3.240	1.392	3.161	8.561	2.345	1.222	1.154	0.243	1.297
BHS	3.300	1.705	3.721	8.701	2.504	1.297	1.732	0.743	1.059
BIH	3.300	2.141	3.927	8.701	2.965	1.316	0.904	0.414	1.295
BLR	3.160	2.086	3.962	8.335	2.933	1.312	0.769	0.757	1.395
BLZ	3.230	1.335	3.926	8.525	2.109	1.313	2.251	0.838	1.902
BOL	3.230	1.340	4.230	8.533	2.114	1.353	5.758	0.860	1.082
BRA	3.240	1.648	3.978	8.558	2.472	1.323	2.251	0.655	1.136
BRB	3.280	1.938	4.339	8.664	2.740	1.359	1.924	0.457	1.499
BWA	3.210	1.251	3.527	8.480	2.053	1.281	6.356	1.231	1.868
CAF	3.230	1.074	3.473	8.522	1.751	1.273	20.472	2.269	2.046
CHE	3.200	2.121	4.389	8.451	2.969	1.370	0.577	0.192	0.815
CHL	3.240	1.747	4.138	8.541	2.570	1.340	1.347	0.386	0.981
CHN	3.300	2.084	3.784	8.716	2.936	1.305	1.347	0.386	0.981
CIV	3.190	1.118	3.349	8.417	1.839	1.259	16.143	2.156	2.522
CMR	3.200	1.103	3.447	8.435	1.841	1.272	14.883	1.780	1.961
COL	3.250	1.581	3.995	8.580	2.378	1.325	2.444	0.627	1.105
CRI	3.230	1.683	4.032	8.533	2.493	1.329	1.347	0.431	0.933
DEU	3.230	2.292	4.472	8.523	3.085	1.372	0.558	0.294	0.778
DNK	3.280	2.034	4.325	8.647	2.848	1.372	0.539	0.254	0.992
DOM	3.230	1.487	4.021	8.523	2.276	1.324	4.332	0.742	1.005
DZA	3.230	1.573	3.924	8.528	2.432	1.318	3.811	0.424	1.202
ECU	3.220	1.458	4.028	8.498	2.254	1.329	3.291	0.516	1.071
EGY	3.150	1.403	3.623	8.322	2.196	1.289	3.291	0.765	1.827
ESP	3.280	2.178	4.522	8.670	3.078	1.380	0.385	0.243	0.859
ETH	3.200	1.087	3.702	8.452	1.802	1.302	9.813	1.065	1.708
FJI	3.300	1.653	3.418	8.705	2.450	1.265	3.349	0.871	1.731
FRA	3.270	2.132	4.652	8.632	2.939	1.393	0.519	0.316	0.793
GBR	3.230	2.017	4.411	8.518	2.836	1.373	0.635	0.294	0.952
GEO	3.190	1.992	4.049	8.422	2.825	1.326	1.809	0.743	1.800
IDN	3.300	1.668	3.604	8.713	2.493	1.285	3.946	0.823	2.012
IND	3.250	1.486	3.635	8.589	2.287	1.288	6.627	0.826	1.530
IRQ	3.270	1.214	3.672	8.642	1.995	1.294	5.296	0.810	1.659
ITA	3.300	2.318	4.549	8.708	3.157	1.382	0.385	0.236	0.908
JAM	3.220	1.585	4.056	8.509	2.379	1.332	2.579	0.602	1.066
JOR	3.330	1.357	3.777	8.796	2.150	1.308	2.713	0.501	1.588
JPN	3.250	2.133	4.842	8.589	3.022	1.421	0.385	0.221	0.779
KAZ	3.160	1.770	3.721	8.336	2.617	1.287	1.809	0.856	1.873
KGZ	3.140	1.530	3.448	8.286	2.342	1.252	2.964	0.757	1.848
KHM	3.300	1.440	3.527	8.706	2.219	1.278	4.736	0.810	1.847
LAO	3.320	1.341	3.547	8.762	2.123	1.279	9.871	0.918	2.000
LBN	3.310	1.647	3.935	8.725	2.447	1.316	1.019	0.423	1.509
LKA	3.240	1.810	3.766	8.540	2.619	1.308	1.539	0.604	1.252
MAR	3.290	1.658	3.705	8.673	2.460	1.295	3.754	0.308	1.751
MDG	3.200	1.166	3.470	8.449	1.895	1.267	7.785	1.041	1.823

(Continued)

TABLE A.5
(CONTINUED)

Country	$\pi_d(1) \times 10^5$	$\pi_d(2) \times 10^3$	$\pi_d(3) \times 10^2$	$\pi_b(1) \times 10^4$	$\pi_b(2) \times 10^2$	$\pi_b(3) \times 10^1$	$\bar{\pi}_n(1) \times 10^5$	$\bar{\pi}_n(2) \times 10^4$	$\bar{\pi}_n(3) \times 10^3$
MEX	3.240	1.526	3.977	8.542	2.331	1.322	2.117	0.581	1.028
MMR	3.350	1.605	3.501	8.855	2.411	1.273	7.804	0.927	1.993
MNG	3.210	1.628	3.517	8.480	2.482	1.271	2.713	1.010	1.682
MOZ	3.190	1.063	3.572	8.430	1.773	1.288	12.288	1.765	1.728
MWI	3.220	1.044	3.561	8.495	1.758	1.287	9.368	1.272	2.079
MYS	3.310	1.555	3.739	8.725	2.371	1.299	1.289	0.550	1.442
NAM	3.190	1.202	3.666	8.433	1.962	1.293	7.186	1.511	1.982
NER	3.160	1.088	3.382	8.346	1.765	1.267	17.714	1.222	2.349
NIC	3.220	1.350	3.876	8.513	2.153	1.308	3.099	0.681	0.991
NOR	3.240	1.935	4.253	8.565	2.782	1.359	0.404	0.215	0.900
NPL	3.250	1.324	3.710	8.589	2.065	1.305	5.257	0.694	1.877
PAK	3.190	1.284	3.704	8.413	2.045	1.300	12.095	0.737	1.811
PAN	3.220	1.567	4.156	8.498	2.379	1.341	2.521	0.497	0.894
PER	3.240	1.566	4.038	8.546	2.397	1.332	2.444	0.564	1.142
PHL	3.320	1.466	3.668	8.751	2.247	1.291	4.332	0.915	1.904
PRY	3.230	1.332	3.931	8.531	2.108	1.321	3.041	0.672	1.198
RWA	3.210	1.195	3.366	8.463	1.927	1.258	6.935	0.942	1.669
SEN	3.190	1.127	3.544	8.417	1.860	1.283	8.306	0.887	2.213
SLE	3.210	1.132	3.492	8.480	1.867	1.277	18.471	2.104	3.109
SLV	3.240	1.424	4.191	8.554	2.195	1.349	2.444	0.855	1.136
TGO	3.210	1.169	3.343	8.462	1.920	1.260	13.256	1.317	2.580
THA	3.350	2.106	4.002	8.840	2.932	1.325	1.924	0.675	1.089
TZA	3.190	1.121	3.432	8.428	1.855	1.270	9.291	1.290	1.808
UGA	3.190	1.008	3.370	8.428	1.709	1.261	9.678	1.456	1.873
USA	3.260	1.917	4.123	8.607	2.710	1.339	1.096	0.517	0.975
VNM	3.280	1.756	3.941	8.669	2.601	1.307	3.156	0.567	0.995
ZAF	3.300	1.459	3.613	8.710	2.287	1.289	6.935	1.535	1.714
ZMB	3.200	1.022	3.473	8.448	1.742	1.273	10.799	1.446	1.790
ZWE	3.200	1.095	3.650	8.454	1.835	1.293	9.678	1.755	1.775

NOTE: This table lists the following calibrated country-specific parameters: the COVID-19 mortality rates ($\pi_d(\cdot)$), the hospitalization probability ($\pi_b(\cdot)$), and the baseline mortality rates ($\bar{\pi}_n(\cdot)$).

TABLE A.6
RESULTS UNDER DIFFERENT ONSET TIMINGS AND LENGTHS OF LOCKDOWNS

	Lower Income			Lower-Middle Income			Upper-Middle Income		
	Mortality Ratio	Δ Children Deaths	Δ Adults Saved	Mortality Ratio	Δ Children Deaths	Δ Adults Saved	Mortality Ratio	Δ Children Deaths	Δ Adults Saved
Baseline	1.758	1.000	1.000	0.585	1.000	1.000	0.061	1.000	1.000
Length + 2w	1.904	1.286	1.190	0.638	1.286	1.171	0.066	1.286	1.185
Length + 4w	2.057	1.572	1.348	0.681	1.572	1.322	0.067	1.569	1.424
Length + 6w	2.190	1.857	1.496	0.700	1.857	1.499	0.062	1.845	1.853
Length + 8w	2.246	2.143	1.684	0.650	2.136	1.795	0.051	2.106	2.729
Start + 2w	0.829	0.991	2.073	0.282	0.994	2.001	0.031	0.986	1.953
Start + 4w	0.417	0.978	4.148	0.148	0.988	3.857	0.017	0.961	3.409
Start + 6w	0.278	0.974	6.401	0.100	0.990	5.535	0.012	0.927	4.830
Start + 8w	0.397	0.992	4.552	0.134	0.999	3.784	0.013	0.930	4.736

Note: The table reports the ratio of children's deaths per adult COVID mortality averted, and the change in the number of children's deaths and adult lives saved when compared to the baseline scenario for alternative reference lockdowns. The middle panel shows the results of prolonging the lockdown by between two and eight weeks. The bottom panel shows the results delaying the lockdown onset by two to eight weeks.

TABLE A.7
RESULTS UNDER DIFFERENT σ 'S

	$\sigma = 5$	$\sigma = 3$ (baseline)	$\sigma = 2.5$ (fertility lit.)	$\sigma = 2$	$\sigma = 1.5$	$\sigma = 1.01$	$\sigma = 1.001$	$\sigma = 1.0001$
	Lower-Income Country Mortality Ratio							
Reference lockdown	1.796	1.758	1.739	1.709	1.664	1.355	1.342	1.343
Optimal lockdown	0.417	0.319	0.288	0.271	0.279	0.232	0.218	0.216
	Δ (No-Lockdown Labor Supply)							
Lower Income	0.007	0.005	0.004	0.003	0.002	0.015	0.037	0.037
Lower-Middle Income	0.009	0.009	0.009	0.008	0.008	0.021	0.039	0.039
Upper-Middle Income	0.013	0.014	0.014	0.015	0.019	0.037	0.045	0.045
High Income	0.026	0.032	0.035	0.041	0.055	0.086	0.088	0.088

NOTE: The top panel reports the number of children's lives lost per adult life saved in Lower-Income Countries under the reference and optimal lockdowns. The bottom panel reports the change in the labor supply without any government-imposed lockdown in each group of countries.

TABLE A.8
BLINDER–OAXACA DECOMPOSITION

	Optimal Lockdown			Criteria-Based Lockdown		
	Overall	Explained	Fraction	Overall	Explained	Fraction
LIC + LMC	0.204 (0.021)			1.061 (0.121)		
UMC + HIC	0.010 (0.002)			0.039 (0.007)		
Difference	0.194 (0.021)			1.021 (0.122)		
Explained	0.283 (0.051)			1.659 (0.208)		
Unexplained	-0.089 (0.042)			-0.638 (0.164)		
Semielasticity		0.226 (0.060)	0.798		1.333 (0.240)	0.803
Population Share, 0-14		0.097 (0.031)	0.343		0.477 (0.172)	0.288
Population Share, 65+		-0.015 (0.023)	-0.053		-0.047 (0.126)	-0.028
Hospital Beds per 1000		-0.002 (0.003)	-0.007		-0.024 (0.020)	-0.014
π_{I1}		-0.001 (0.011)	-0.004		0.032 (0.038)	0.019
π_{I2}		0.009 (0.006)	0.032		0.061 (0.034)	0.037
π_{I3}		-0.033 (0.014)	-0.117		-0.173 (0.059)	-0.104
Constant						
N	85			85		

NOTE: This table reports the two-way Blinder–Oaxaca decomposition of the expected number of children lives lost per COVID-19 fatality averted by country groups. The first group is the low-income countries (LIC) and the lower-middle-income countries (LMC); the second group is the upper-middle-income countries (UMC) and the high-income countries (HIC).

A.4 Blinder–Oaxaca Decomposition. One way to assess the relative importance of the various mechanisms in determining the expected number of child lives lost per COVID-19 fatality averted is to split all sample countries according to a binary distinction of low- versus high-income, and then conduct an Oaxaca–Blinder decomposition including all factors explored individually above. We include the semielasticity of child mortality with respect to income ($\nu(1)$), the population age distribution (ℓ_1 and ℓ_3), hospital capacity (κ), and the calibrated transmission parameters (π_{I1} , π_{I2} , and π_{I3}) on the right-hand side (RHS) of the decomposition. Table A.8 reports the results of the decomposition exercise.

The columns “overall” reports the differences in means across country groups and the overall explanatory power of the RHS variables. Three mechanisms absorb most of the cross-country variation in the number of children’s lives lost per COVID-19 fatality averted (these findings hold for the reference lockdown as well as the optimal lockdown discussed below). The most influential factor is the semielasticity of child mortality to the economic contraction, which accounts for 80% of the explained variation across the two country groups. The share of the population under 15 years of age is also a significant factor, accounting for 29% of the variation. These two factors “explain” more than 100% of the variation across countries since other considered factors decrease the cross-country variation in the ratio of child deaths to averted COVID-19 deaths. The most significant factor that decreases this variation is the community transmission parameter. When community-based transmission constitutes a larger share of total COVID-19 disease transmission, any lockdown policy will be less effective in averting COVID-19 mortality, resulting in higher ratios of child deaths to COVID-19

deaths averted. These three factors are the only significant factors in the decomposition exercise. Other factors such as the share of the population 60 years or older, hospital capacity, or the work- and consumption-related transmission parameters are not especially influential in explaining the observed variation in the trade-off between child- and COVID-19-mortality.

REFERENCES

- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON, "Optimal Targeted Lockdowns in a Multi-Group SIR Model," *American Economic Review: Insights* 3 (2021), 487–502.
- AHMED, T., T. ROBERTON, T. MOEHS, J. ALFRED, M. BAYE, M. DIABATE, H. KIARIE, P. MBAKA, N. MOHAMUD, C. MWANSAMBO, Y. NDIAYE, C. NZELU, A. OFOSU, T. RAHARISON, H. SADAT, F. A. SMART, H. UDDIN, N. WENDRAD, C. S. WESSEH, M. L. YANSANE, S. YUMA, P. VERGEER, P. M. HANSEN, J. A. FRIEDMAN, AND G. SHAPIRA, "Indirect Effects on Maternal and Child Mortality from the COVID-19 Pandemic: Evidence from Disruptions in Healthcare Utilization in 18 Low-and Middle-Income Countries," *SSRN Electronic Journal*, 2021.
- ALON, T. M., M. KIM, D. LAGAKOS, AND M. VANVUREN, "How Should Policy Responses to the COVID-19 Pandemic Differ in the Developing World?," NBER Working Paper No. 27273, May 2020.
- ALVAREZ, F., D. ARGENTE, AND F. LIPPI, "A Simple Planning Problem for COVID-19 Lock-down, Testing, and Tracing," *American Economic Review: Insights* 3 (September 2021), 367–82.
- AMINJONOV, U., O. BARGAIN, AND T. BERNARD, "Gimme Shelter. Social Distancing and Income Support in times of Pandemic," Technical Report, Groupe de Recherche en Economie Théorique et Appliquée (GREThA), 2021.
- ANDREASEN, V., C. VIBOUD, AND L. SIMONSEN, "Epidemiologic Characterization of the 1918 Influenza Pandemic Summer Wave in Copenhagen: Implications for Pandemic Control Strategies," *The Journal of Infectious Diseases* 197 (2008), 270–8.
- ATKESON, A., "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios," NBER Working Paper 26867, March 2020.
- , K. KOPECKY, AND T. ZHA, "Four Stylized Facts about COVID-19," NBER Working Paper 27719, August 2020.
- BAIRD, S., J. FRIEDMAN, AND N. SCHADY, "Aggregate Income Shocks and Infant Mortality in the Developing World," *Review of Economics and Statistics* 93 (2011), 847–56.
- , D. MCKENZIE, AND B. ÖZLER, "The Effects of Cash Transfers on Adult Labor Market Outcomes," *IZA Journal of Development and Migration* 8 (2018), 1–20.
- , AND E. FARHI, "Nonlinear Production Networks with an Application to the Covid-19 Crisis," Mimeo, UCLA and Harvard, April 2020.
- BAQAEI, D., AND E. FARHI, "Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis," *American Economic Review* (February 2021).
- BARNETT-HOWELL, Z., AND M. MOBARAK, "The Benefits and Costs of Social Distancing in Rich and Poor Countries," arXiv:2004.04867, April 2020.
- BARROT, J. N., B. GRASSI, AND J. SAUVAGNAT, "Sectoral Effects of Social Distancing," *American Economic Association Papers and Proceedings* 111 (May 2021), 277–81.
- BEYER, R., T. JAIN, AND S. SINHA, "Lights Out? COVID-19 Containment Policies and Economic Activity," Available at SSRN: <https://ssrn.com/abstract=3856050> or <https://doi.org/10.2139/ssrn.3856050>, May 2021.
- BHALOTRA, S., "Fatal Fluctuations? Cyclicalities in Infant Mortality in India," *Journal of Development Economics* 93 (2010), 7–19.
- BONADIO, B., Z. HUO, A. A. LEVCHENKO, AND N. PANDALAI-NAYAR, "Global Supply Chains in the Pandemic," *Journal of International Economics* 133 (November 2021), 103534.
- CRUCES, G., P. GLÜZMANN, AND L. F. L. CALVA, "Economic Crises, Maternal and Infant Mortality, Low Birth Weight and Enrollment Rates: Evidence from Argentina's Downturns," *World Development* 40 (2012), 303–14.
- CUTLER, D. M., F. KNAUL, R. LOZANO, O. MÉNDEZ, AND B. ZURITA, "Financial Crisis, Health Outcomes and Ageing: Mexico in the 1980s and 1990s," *Journal of Public Economics* 84 (2002), 279–303.
- DEHEJIA, R., AND A. LLERAS-MUNEY, "Booms, Busts, and Babies' Health," *The Quarterly Journal of Economics* 119 (2004), 1091–130.
- DOEPKE, M., "Accounting for Fertility Decline During the Transition to Growth," *Journal of Economic Growth* 9 (2004), 347–83.
- EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT, "The Macroeconomics of Epidemics," *Review of Financial Studies* 34 (November 2021), 5149–87.
- FERGUSON, N. M., D. A. CUMMINGS, C. FRASER, J. C. CAJKA, P. C. COOLEY, AND D. S. BURKE, "Strategies for Mitigating an Influenza Pandemic," *Nature* 442 (2006), 448–52.

- FERNÁNDEZ-VILLAVARDE, J., AND C. I. JONES, "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities," Working Paper 27128, National Bureau of Economic Research, May 2020.
- FERREIRA, F., AND N. SCHADY, "Aggregate Economic Shocks, Child Schooling, and Child Health," *World Bank Research Observer* 24 (August 2009), 147–81.
- FLAXMAN, S., S. MISHRA, A. GANDY, H. J. T. UNWIN, T. A. MELLAN, H. COUPLAND, C. WHITTAKER, H. ZHU, T. BERAH, J. W. EATON, AND M. MONOD, "Estimating the Effects of Non-Pharmaceutical Interventions on COVID-19 in Europe," *Nature* 584 (2020), 257–61.
- FRIEDMAN, J., AND N. SCHADY, "How many Infants Likely Died in Africa as a Result of the 2008–2009 Global Financial Crisis?," *Health Economics* 22 (2013), 611–22.
- GARSKE, T., A. CORI, A. ARIYARAJAH, I. M. BLAKE, I. DORIGATTI, T. ECKMANN, C. FRASER, W. HINSLEY, T. JOMBART, H. L. MILLS, AND G. NEDJATI-GILANI, "Heterogeneities in the Case Fatality Ratio in the West African Ebola Outbreak 2013–2016," *Philosophical Transactions of the Royal Society B: Biological Sciences* 372 (2017), 20160308.
- GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RÍOS-RULL, "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic," CEPR Discussion Paper 14606, April 2020.
- GREENWOOD, J., Z. HERCOWITZ, AND G. W. HUFFMAN, "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review* 78 (June 1988), 402–17.
- HALL, R. E., C. I. JONES, AND P. J. KLENOW, "Trading off Consumption and Covid-19 Deaths," *Federal Reserve Bank of Minneapolis Quarterly Review* 42 (June 2020).
- JIA, N., D. FENG, L.-Q. FANG, J. H. RICHARDUS, X.-N. HAN, W.-C. CAO, AND S. J. DE VLAS, "Case Fatality of SARS in Mainland China and Associated Risk Factors," *Tropical Medicine & International Health* 14 (2009), 21–7.
- JOHNSTONE-ROBERTSON, S. P., D. MARK, C. MORROW, K. MIDDELKOOP, M. CHISWELL, L. D. AQUINO, L.-G. BEKKER, AND R. WOOD, "Social Mixing Patterns within a South African Township Community: Implications for Respiratory Disease Transmission and Control," *American journal of Epidemiology* 174 (2011), 1246–55.
- KAPLAN, G., B. MOLL, AND G. VIOLANTE, "The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the U.S.," Mimeo, Chicago, LSE, and Princeton, September 2020.
- KARLBERG, J., D. CHONG, AND W. LAI, "Do Men have a Higher Case Fatality Rate of Severe Acute Respiratory Syndrome than Women Do?," *American Journal of Epidemiology* 159 (2004), 229–31.
- KIM, Y. E., AND N. V. LOAYZA, "Economic Loss from COVID-19 Fatalities Across Countries: A VSL Approach," *Applied Economics Letters* 7 (2021), 644–50.
- KRUEGER, D., H. UHLIG, AND T. XIE, "Macroeconomic Dynamics and Reallocation in an Epidemic: Evaluating the 'Swedish Solution'," *Economic Policy* (May 2022), <https://doi.org/10.1093/epolic/eiac010>.
- LAVEZZO, E., E. FRANCHIN, C. CIAVARELLA, G. CUOMO-DANNENBURG, L. BARZON, C. DEL VECCHIO, L. ROSSI, R. MANGANELLI, A. LOREGIAN, N. NAVARIN, AND D. ABATE, "Suppression of a SARS-CoV-2 Outbreak in the Italian Municipality of Vo'," *Nature* 584 (2020), 425–9.
- LEE, B. Y., S. T. BROWN, P. C. COOLEY, R. K. ZIMMERMAN, W. D. WHEATON, S. M. ZIMMER, J. J. GREFENSTETTE, T.-M. ASSI, T. J. FURPHY, D. K. WAGENER, AND D. S. BURKE, "A Computer Simulation of Employee Vaccination to Mitigate an Influenza Epidemic," *American Journal of Preventive Medicine* 38 (2010), 247–57.
- LEI, H., X. XU, S. XIAO, X. WU, AND Y. SHU, "Household Transmission of COVID-19—a Systematic Review and Meta-Analysis," *Journal of Infection* 81 (2020), 979–97.
- LEVIN, A. T., W. P. HANAGE, N. OWUSU-BOATEY, K. B. COCHRAN, S. P. WALSH, AND G. MEYEROWITZ-KATZ, "Assessing the Age Specificity of Infection Fatality Rates for Covid-19: Systematic Review, Meta-Analysis, and Public Policy Implications," *European journal of epidemiology* 35 (2020), 1123–38.
- LI, R., S. PEI, B. CHEN, Y. SONG, T. ZHANG, W. YANG, AND J. SHAMAN, "Substantial Undocumented Infection Facilitates the Rapid Dissemination of Novel Coronavirus (SARS-CoV-2)," *Science* 368 (2020), 489–93.
- LOAYZA, N., "Macroeconomic Dynamics and Reallocation in an Epidemic," World Bank Research and Policy Brief No. 35, May 2020.
- MEYEROWITZ-KATZ, G., AND L. MERONE, "A Systematic Review and Meta-Analysis of Published Research Data on COVID-19 Infection-Fatality Rates," *International Journal of Infectious Diseases* 101 (2020), 138–48.
- OLSON, D. R., L. SIMONSEN, P. J. EDELSON, AND S. S. MORSE, "Epidemiological Evidence of an Early Wave of the 1918 Influenza Pandemic in New York City," *Proceedings of the National Academy of Sciences* 102 (2005), 11059–63.
- PRITCHETT, L., AND L. H. SUMMERS, "Wealthier is Healthier," *Journal of Human Resources* 31 (1996), 841.
- RAVALLION, M., "Pandemic Policies in Poor Places," CGD Note, April 2020.
- ROBERTSON, T., E. D. CARTER, V. B. CHOU, A. R. STEGMULLER, B. D. JACKSON, Y. TAM, T. SAWADOGO-LEWIS, AND N. WALKER, "Early Estimates of the Indirect Effects of the COVID-19 Pandemic On Mater-

- nal and Child Mortality in Low-Income and Middle-Income Countries: A Modelling Study,” *The Lancet Global Health* 8 (2020), e901–8.
- RUHM, C. J., “Are Recessions Good for Your Health?,” *The Quarterly Journal of Economics* 115 (2000), 617–50.
- SINGER, P., AND M. PLANT, “When Will the Pandemic Cure Be Worse Than the Disease?,” Project Syndicate, April 2020.
- SLOVIC, P., AND E. PETERS, “Risk Perception and Affect,” *Current Directions in Psychological Science* 15 (2006), 322–5.
- STEVENS, A. H., D. L. MILLER, M. E. PAGE, AND M. FILIPSKI, “The Best of Times, the Worst of Times: Understanding Pro-Cyclical Mortality,” *American Economic Journal: Economic Policy* 7 (2015), 279–311.
- STRINGHINI, S., A. WISNIAK, G. PIUMATTI, A. S. AZMAN, S. A. LAUER, H. BAYSSON, D. DE RIDDER, D. PETROVIC, S. SCHREMPFT, K. MARCUS, AND S. YERLY, “Repeated Seroprevalence of Anti-SARS-CoV-2 IgG Antibodies in a Population-Based Sample from Geneva, Switzerland,” *medRxiv* (2020).
- VERITY, R., L. C. OKELL, I. DORIGATTI, P. WINSKILL, C. WHITTAKER, N. IMAI, G. CUOMO-DANNENBURG, H. THOMPSON, P. G. WALKER, H. FU, et al., “Estimates of the Severity of Coronavirus Disease 2019: A Model-Based Analysis,” *The Lancet Infectious Diseases* 20 (2020), 669–77.
- VISCUSI, W. K., AND J. E. ALDY, “The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World,” *Journal of Risk and Uncertainty* 27 (2003), 5–76.
- WALKER, P. G. T., C. WHITTAKER, AND O. WATSON, “The Global Impact of COVID-19 and Strategies for Mitigation and Suppression: WHO Collaborating Centre for Infectious Disease Modelling,” MRC Centre for Global Infectious Disease Analysis, Abdul Latif Jameel Institute for Disease and Emergency Analytics, Imperial College London (2020).
- YANG, S., P. CAO, P. DU, Z. WU, Z. ZHUANG, L. YANG, X. YU, Q. ZHOU, X. FENG, X. WANG, AND W. LI, “Early Estimation of the Case Fatality Rate of Covid-19 in Mainland China: A Data-Driven Analysis,” *Annals of Translational Medicine* 8 (2020), 128.