

The Macroeconomics of Pandemics in Developing Countries: an Application to Uganda

Tillmann von Carnap^a

Ingvild Almås^b

Tessa Bold^c

Selene Ghisolfi^d

Justin Sandefur^e

Abstract

How should policies to control the SARS-CoV-2 pandemic differ across countries? We extend recent contributions integrating economic and epidemiological models for the United States to a developing country context, Uganda. Differences in demography, comorbidities, and health systems affect the mortality risk; lower incomes affect agents' willingness to forego consumption to reduce disease risk. We find that, for a broad range of life valuations supported by the literature, optimal containment is significantly less restrictive in the latter context, a normative picture that contradicts the positive of similarly strict measures. We explore biased beliefs on the disease's risks as a possible explanation.

^aPhD Student, Institute for International Economic Studies, Stockholm University, Corresponding author: tillmann.voncarnap@iies.su.se

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

^cAssociate Professor, Institute for International Economic Studies, Stockholm University. tessa.bold@iies.su.se

^dPhD Student, Institute for International Economic Studies, Stockholm University and LEAP, Bocconi University, Milan, selene.ghisolfi@iies.su.se

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

Acknowledgements: We thank Timo Boppart, Tobias Broer, Raymond Guiteras, Per Krusell, David Lagakos, Mahreen Mahmud, Mushfiq Mobarak, Rohini Pande, Debraj Ray, Emma Riley, Jakob Svensson, Lore Vandewalle, William Kip Viscusi, Has van Vlokhoven and seminar participants at the IIES Brown Bag and the TARC workshop on Covid-19 and Developing Countries for helpful comments. An earlier version of this paper has been published in the 27th issue of the CEPR Covid Economics series.

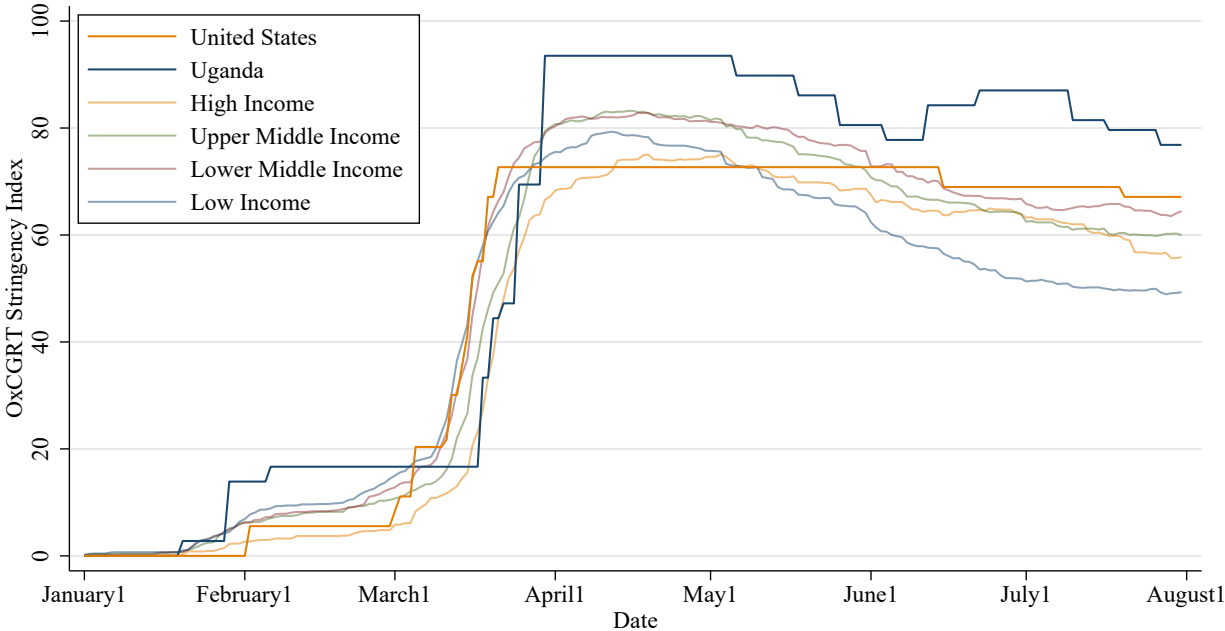
1 Introduction

The ongoing COVID-19 pandemic has led governments around the world to impose unprecedented restrictions on economic activity, with surprising initial uniformity across countries at all income levels (Figure 1). In the United States, a survey in late March 2020 found zero leading economists disagreed that the policy response to the pandemic should involve “a very large contraction in economic activity until the spread of infections has dropped significantly” (IGM Forum, 2020). For the developing world, however, economists have expressed reservations about similar policy prescriptions (Ray et al., 2020; Ray and Subramanian, 2020; Barnett-Howell and Mobarak, 2020; Ravallion, 2020).

In this paper, we study the rationale for possible policy responses by exploring two factors that are crucial in order to understand differences across developing and developed countries: i) the mortality risk as measured through the infection fatality rate (IFR) - which incorporates the effects of demography, comorbidities, and health system capacity, ii) and poverty - which increases the utility cost of the lockdown. We also study how the value that a social planner’s attaches to saving a life matters for optimal policies.

More specifically, we extend recent work integrating economic behavior into an epidemiological Susceptible - Infected - Recovered (SIR) model to take account of subsistence constraints, demography-dependent fatality rates and context-specific valuations of a statistical life. We use our model to compare welfare-optimizing policies across contexts that differ in terms of both income and demography, presenting the

Figure 1: OxCGRT Stringency Index in 2020 - Uganda, United States and average per income group



The figure shows the average level of restrictions over time in the four income groups defined by the World Bank and our study countries as measured by the Oxford Coronavirus Government Response Stringency Index (Hale et al., 2020).

examples of the United States and Uganda. We focus on the latter as an example of a developing country highlighting the two salient features we explore.

First, Uganda is the third youngest country in the world, its median age - 16 - being less than half that of the United States - 38 - implying the probability of death from the disease may be considerably lower. In our framework, this risk differentially affects agents' labor supply and consumption decisions and, in turn, optimal policy. A key statistic summarizing these demographic differences is the IFR, the share of infected people dying. In Ghisolfi et al. (2020), we show that a population's age and comorbidity structure predicts wide variation in this statistic across contexts, and that even after adjusting for the lower capacity of Uganda's health system relative to the United States, the average infected person may be less than half as likely to die from COVID-19 in Uganda compared to the United States (0.33% vs. 0.79%).¹ We note that recent evidence points to the possibility of an even larger differential in the IFR between rich and poor countries (Nordling, 2020).

Second, Uganda's GDP per capita - \$710 - is substantially lower than that of the United States - \$56,000 - and very close to the average in low-income countries. At lower incomes, foregone consumption due to pandemic control (both voluntary and policy-induced) implies a larger welfare loss, affecting agents' and policymakers' optimal choices.

Despite these differences in fatality risk and incomes, Uganda, as most low-income countries, has responded to the virus relatively early, strictly and sustained (Figure 1). Available figures suggest successful initial containment: by early August, there had been about 1,200 confirmed cases and less than ten reported deaths.² There is however emerging evidence of widespread economic hardship as a result of the lockdown policies: Mahmud and Riley (2020) report that households from a rural sample had seen their incomes decrease by 60%, reduced their food expenditure by 40%, and worked substantially more on their own farms and subsequently less in other sectors.

Our study shows how the lower fatality risk changes the implicit valuation of the utility of the living against the number of deaths. Furthermore, at lower income levels and close to subsistence, consumption adjustments are less elastic, and hence less spontaneous adjustment can be expected. This is in line with evidence from the United States and India: Chetty et al. (2020) report that half of the observed reduction in consumer spending during the early stages of the United States outbreak came from households in the top quartile of the income distribution, particularly in sectors requiring physical interactions. In Delhi, occupation groups with higher incomes tend to cite fear of the virus as the main reason for not returning to

¹We construct IFRs starting from available estimates on the number of comorbidities among COVID-19 deaths in Italy and the gender and age structure of fatalities in France. We then adjust these figures to take account of the age and comorbidity distribution of all countries, as well as the health system's capacity to deal with pulmonary diseases.

²Source: ourworldindata.org

work after the relaxation of the initially strict lockdown was scaled back, while poorer households are mostly affected through the lack of job opportunities (Desai and Pramanik, 2020). This suggests richer households can afford taking precautions that may be economically challenging for poorer ones.

We add a developing country’s perspective to recent work incorporating economic decision-making into the SIR framework. According to these integrated models, agents facing contagion risk will voluntarily reduce their economic activity, thus partly containing the spread of the epidemic (Toxvaerd, 2020; Garibaldi et al., 2020; Chudik et al., 2020). However, analyses suggest that, from a social welfare perspective, further government action is justified by agents’ failure to internalize their own contribution to the spread of the epidemic (Eichenbaum et al., 2020; Farboodi et al., 2020; Krueger et al., 2020; Glover et al., 2020; Alvarez et al., 2020). Few model-based studies have focused on developing countries and the efficacy of the aforementioned mechanisms in their contexts, Alon et al. (2020) and Hausmann and Schetter (2020) being exceptions we discuss below.

For our exercise, we start from the model by Eichenbaum et al. (2020, henceforth ERT), in which agents expose themselves to infection risk when working and consuming, and, realizing this danger, reduce economic activity as infection risk rises. However, they do not weigh the impact of their own labor and consumption decisions on the pandemic’s spread, creating an externality the social planner seeks to internalize. The optimal policy maximizes the present value of societal utility, taking the perspective of a social planner at the beginning of the pandemic. Components which the social planner balances are i) the utility the agents derive from consumption and leisure, ii) the disutility of foregone consumption from lower productivity when infected, and iii) the disutility of dying.

In line with other studies, we confirm that, if a vaccine is unlikely to become available and a survived infection gives immunity, the ‘optimal containment rate’ follows the share of infected in the population, i.e. it discourages economic activity more when the risk of contracting or spreading the disease is higher, thus guiding the population to herd immunity. Although never restricting economic activity totally, our United States calibration yields a strong and sustained discouragement of consumption, equivalent to a 55% consumption tax over the first year of the epidemic.

A central assumption of the ERT model is that the disutility of death is equal to the foregone utility of living. Placing a monetary value on a life is challenging, and some would even argue unethical. However, decision makers in all societies are making trade-offs that, implicitly or explicitly, assign monetary values to lives. Examples are implicit valuations in budget posts for public health services and explicit valuations when using calculated risks of deaths in cost-benefit analysis of infrastructure projects. In the United States, a common choice for the value of a statistical life (VSL) is around \$10 million per life, which is close to the value we inherit from ERT and use in the United States calibration. However, the monetary value of

a life will in reality depend also on the decision maker's budget and as such valuations can be expected to be higher in richer than in poorer countries. There neither exists a broadly accepted way of scaling the consensus based evaluations to contexts with lower incomes, nor a collection of reference values used in poorer countries (Viscusi and Masterman, 2017).

We review the literature estimating VSLs in developing countries, and find that scaling rich countries' VSLs suggests values up to 100 times higher than those found in micro-studies from the experimental development economics literature. As our model takes the social planner's perspective, and the United States consensus-based value reflects cost-benefit analyses used for policy rather than aggregated individual willingness to pay (WTP), we purposefully target a valuation of life in Uganda to reflect some value that we observe in policy makers' choices. One such value is derived from the spending prioritization of governments on keeping their citizens healthy. In particular, we adjust the consensus value for the United States with the ratio between the Ugandan and American per-capita health spending. This value lands at about \$31,000, a quarter of what a simple scaling by GDP per capita would return, but at least ten times higher than results from the experimental literature and three times higher than what we get by calibrating our model to Uganda without targeting any value, thereby sticking closest to the social planner's objective function in the United States calibration.

We find that the difference in terms of mortality between a simple SIR model and one integrating welfare optimizing policy is much smaller in Uganda than in the United States. While in the latter the optimal containment policy reduces mortality by 34% relative to the pure epidemiological model without behavioral adjustments, in Uganda the reduction is only 10%. Both of our suggested factors contribute about equally to this result: first, lower IFRs in developing countries make for a lower aggregate disutility of infection. Second, even in the face of contagion risk and containment measures, a poorer agent experiences larger relative utility losses when reducing consumption. The former makes the behavioral responses to any tax less efficient from a health perspective, and the latter makes the social planner less willing to introduce a tax that reduces consumption. Thus, both forces suggest that optimal policies should be less restrictive in developing countries (characterized by lower incomes and younger populations with lower predicted IFRs) than in richer economies. The difference increases when valuing deaths at what individuals in developing countries reveal through health prevention behavior, and remains, albeit more muted, when applying the upper end of values found in the literature using cross-country estimations.

We further extend our model to explain a seemingly paradoxical observation from recent surveys in low-income countries: respondents state high rates of agreement with - according to our analysis overly strict - lockdowns imposed by governments, while at the same time experiencing large income losses (Moscoviz and Le Nestour, 2020). If agents perceive an exaggerated risk of contracting the disease and/or of dying from

it, their voluntary adjustment to protect themselves may coincide with the effects of the strict measures introduced by governments. Unless they update their beliefs, agents may well agree with a painfully strict lockdown.

Our study complements independent research by Alon et al. (2020) and Hausmann and Schetter (2020), whose heterogeneous agents models include various characteristics of developing country state and health system capacity to assess how different restrictions affect welfare in rich and poor countries. As the former model assumes that workers in the informal sector cannot be shielded from the disease by a lockdown, the authors argue that containment policy will be less effective in countries with larger informal sectors. Our approach highlights a similar effect, though grounded in the utility maximization of agents and affecting the optimal policy of the social planner: when faced with a given risk of contracting the disease through economic activity, poorer agents will rationally reduce their exposure less, requiring relatively stricter policies to achieve the same reduction in deaths. Alon et al. (2020) further emphasize that demographic differences, as captured by the country-varying IFR in our framework, account for most of the differences in mortality rates between their modeled rich and poor countries. In contrast, our approach highlights that lower mortality risk may not only mechanically affect the overall death rate, but also individual-level optimization and adaptation behavior. In line with their findings, we show that a given containment policy is about half as efficient in averting deaths in Uganda per unit of GDP loss than in the United States. Finally, in contrast to Hausmann and Schetter (2020)'s two-period model of households facing subsistence constraints, we jointly model the full paths of infections and optimal policy.

In the following section, we lay out existing VSL estimates for poorer countries. We then present our epidemiological and economic framework and their calibrations. In Section 4, we present the results before concluding.

2 The value of a statistical life in developing countries

A critical choice in our and other studies of optimal policy in a pandemic is the social planner's willingness to pay to reduce the expected number of fatalities by one, or equivalently, the cost assigned by society to one preventable death (Acemoglu et al., 2020). This 'value of a statistical life' is the subject of substantial discussion in both academic and policy circles (Adler, 2020), and a range of values has been adopted by decision makers (OECD, 2012). Estimations by economists are usually based on compensating variations of incomes for accepting a higher death risk, elicited either through hypothetical choice experiments or from wage differentials in jobs varying by fatality risk. Both approaches have been subject to critique, and the evidence is focused on developed countries. Aiming to fill this gap, Viscusi and Masterman (2017) collect

VSL estimates from contexts varying by income and estimate the cross-country income elasticity of the VSL. With their central estimate of one, they calculate a VSL of \$120,000 for Uganda, slightly higher than what we arrive at - \$86,000 - if using unit elasticity but scaling our reference VSL for the United States with the ratio of our targeted incomes.

In contrast, a review by the OECD (2012) recommends the use of an income elasticity of 0.8, yielding a VSL for Uganda of \$218,000, again using the ratio of incomes used in our calibration (see Table 1a). It is worth noting that even with similar methodologies, we arrive at quite different evaluations – the examples here vary by a factor of almost three.

The values underlying the analyses in Viscusi and Masterman (2017) and OECD (2012) are mostly derived from wage differentials, and stand in stark contrast to the few studies estimating VSLs in poorer countries from revealed preferences (for a recent review, see Robinson et al. (2019)). Kremer et al. (2011) and Ito and Zhang (2020), respectively, observe people’s willingness to travel to safe water sources in rural Kenya and to install air purifiers in China, and translate these to WTP for reducing or eliminating death risk of children from diarrhea and air pollution. Berry et al. (2020) record peoples’ valuations of water filters directly using incentive compatible elicitation methods. These can be translated into VSLs using their estimates of the filters’ effectiveness. Scaled by the respective values of GDP per capita to Uganda with unit elasticity, these studies’ estimates translate into VSLs of between \$600 and \$3,400.

Given the small number of available estimates, we also derive a value from the WTP for insecticide-treated bednets in Cohen and Dupas (2010). The paper states a reduction in under-5 mortality risk of 18%, reducing the death risk for children alive in 2007 from 6.91% to 5.67%. Hence, for avoiding one statistical death, $1/(6.91-5.67)=80.4$ bednets would have to be purchased. Kenyan households’ stated WTP of US\$1.71 for one bednet implies a valuation of a statistical death of $80.4*US\$1.71 = US\137.5 , translating into US\$105 in Uganda.

As evident from Table 1b, VSL estimates based on revealed preferences for preventive health behavior are orders of magnitude lower than those based on scaling estimates from rich countries by a country’s income. A small part of the difference may come from study samples on average being poorer than the overall population. Counteracting this selection problem is the fact that the values are estimated for individuals much younger than those at highest risk from COVID-19. We interpret the low revealed VSLs as reflecting the multiple dimensions in which poor households have to manage fatality risk, not just from diseases but also from low incomes more generally.

While we note the ethical and methodological challenges associated with scaling a small, perhaps inadequately understood death probability from a specific cause to the deterministic outcome of one death from COVID-19, we also emphasize that any model illustrating, and any policymaker contemplating, trade-offs

Table 1: VSL estimates for Uganda

(a) from cross-country transfer calculations				
Source			Income elasticity	VSL scaled to Ug.
<i>Published estimates</i>				
Viscusi & Masterman, 2017			1	US\$ 120,000
<i>Derived estimates</i>				
Viscusi & Masterman, 2017			1	US\$ 86,000
OECD, 2012			0.8	US\$ 218,000
(b) from revealed preference experimental studies				
Source	Year	Context	Estimated VSL	VSL scaled to Ug.
<i>Published estimates</i>				
Kremer et al., 2011	2007	rural Kenya	US\$ 769 - 2,715	US\$ 594 - 2,100
Berry et al., 2020	2009-10	rural Ghana	US\$ 3,604	US\$ 1,830
Ito & Zhang, 2020	2006-14	urban China	US\$ 34,580	US\$ 3,386
<i>Derived estimates</i>				
Cohen & Dupas, 2010	2007	rural Kenya	US\$ 137	US\$ 105

Derived estimates in panel (b) take as inputs stated income elasticities, calibration targets for income from Section 3.1 and United States VSL from ERT. Values of statistical life year in Ito and Zhang (2020) scaled to VSLs by multiplying with Chinese life expectancy. Values scaled to Uganda by ratio between respective country's and Uganda's GDP per capita.

between reducing death risk and other sources of disutility must assign, implicitly or explicitly, a value, often based on a small evidence base. A common criticism of VSL estimates based on WTP is that due to liquidity constraints, subjects may not state a 'true' valuation, but one that reflects limited spending possibilities on health protection behavior. This concern looms larger when comparing estimates across countries than it does here, where we are illustrating a situation in which individuals manage disease risks from multiple sources while being affected by a policy aimed at reducing only one specific risk. Actual lockdowns in developing countries were often not implemented in combination with extended social safety nets, thus also increasing households' exposure to non-COVID-19 fatality risks (see Ray and Subramanian (2020) for the case of India).

Our benchmark value for the Ugandan VSL reflects the prioritization taken by governments, not necessarily by individuals, and at \$31,000 lies in between the two sets of VSLs in the existing literature:

$$VSL_{UG} = VSL_{US} * \frac{Health\ Spending_{UG}/Pop_{UG}}{Health\ Spending_{US}/Pop_{US}}$$

We explore in Section 4.5 how this critical parameter affects the optimal policy. For the United States, we follow ERT's calibration to \$9.3 million, a value that is validated by its use in United States public agency deliberations.

3 Model

Our model combines rationally optimizing agents and an epidemiological model, where the share of the population that is currently either susceptible to, infected with or recovered from a disease, evolves according to a set of parameters. The epidemic starts with an exogenous share of the population being infected. A parameter β denotes the rate at which infected people contact susceptibles and transmit the virus. Once infected, individuals recover at rate γ and are thereafter assumed to be immune. Depending on the relative size of β and γ , infections die out quickly (if people recover at a higher rate than infecting new ones), or the number of infected rises exponentially, until there are only few susceptibles left. At that stage, a sufficiently large share of the population has acquired immunity to the disease and exogenous new infections will not cause a new epidemic. In this model, deaths are recorded as a simple share of those contracting the disease. A caveat with analyses using the SIR model is that its parameters are still surrounded by significant uncertainty (Avery et al., 2020).

The key variable linking the economy and the epidemic is the number of newly infected people in a given period, denoted by T_t :

$$T_t = \pi_{s1} (S_t C_t^s) (I_t C_t^i) + \pi_{s2} (S_t N_t^s) (I_t N_t^i) + \pi_{s3} S_t I_t$$

The first two terms on the right denote new infections coming from interactions while consuming, or more broadly defined, spending money, and working. It is higher the more susceptible (S_t) and infected (I_t) people there are, and the more each group consumes and works ($C_t^s, C_t^i, N_t^s, N_t^i$). The third term captures infections from random interactions outside of work or consumption activities. The π_s parameters govern the likelihood of getting infected from either source. Their estimation is key for the way the epidemic plays out. In line with ERT, we assume that evidence from other epidemics also applies to the current one, in that 1/6 of infections take place at work, 1/6 while consuming, and 2/3 from random interactions. Furthermore we calibrate the π_s using steady state values of hours worked and consumption, targeting a final epidemic size of 60%.

We note that we so far lack the data to validate the transmission parameters of our epidemiological model. In particular, there are neither broad studies of the loci of new infections in poorer settings (during economic activity or within the household), nor of the likelihood of a contact with an infected person resulting in transmission. Measures of these two quantities would affect the transmission probabilities calibrated from the procedure outlined above. If transmissions from consumption and work are less important, and perceived so, in developing countries, the behavioral adjustments in our model would be weaker, and vice versa.

Absent better evidence, we take on board ERT's assumptions about the basic patterns of the pandemic,

implying a basic reproductive number R_0 of 1.5. This is at the lower end of estimates reported for the early stage of the pandemic in Liu et al. (2020).³ We keep the transmission probabilities constant with respect to ERT, allowing for comparability along that dimension. We see this as a conservative way to model the epidemic in developing countries compared to richer settings, as infected individuals there may be more likely to interact with susceptibles outdoors, and lower mobility should make local outbreaks slower to spread across regions.

In the model's economic part, a continuum of representative agents are trading off consumption c and labor n :

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

Agents cannot consume less than the subsistence level \bar{c} and face a budget constraint linking government action to individual consumption through a discouragement on consumption μ_{ct} (the 'containment rate'), the proceeds of which are immediately rebated as a lump sum Γ_t :

$$(1 + \mu_{ct})c_t = w_t n_t + \Gamma_t$$

θ and the wage rate are calculated from the model's steady state, targeting the average weekly income and the number of hours worked.

The chosen subsistence level affects how much agents adjust economic activity to infection risk. We are not aware of nationally representative data from Uganda or a comparable context tracing out this response. The intuition of our calibration, that poorer agents will be able to adjust less, is supported by evidence from the United States, showing that richer households stood for the bulk of spending reductions during the early stages of the pandemic (Chetty et al., 2020).

Agents maximize their lifetime utility, components of which depend on their current infection status.

$$\begin{aligned} \text{Susceptible:} & \quad U_t^s = u(c_t^s, n_t^s) + \beta [(1 - \tau_t) U_{t+1}^s + \tau_t U_{t+1}^i] \\ \text{where:} & \quad \tau_t = \pi_{s1} c_t^s (I_t C_t^I) + \pi_{s2} n_t^s (I_t N_t^I) + \pi_{s3} I_t \\ \text{Infected:} & \quad U_t^i = u(c_t^i, n_t^i) + \beta [(1 - \pi_r - \pi_d) U_{t+1}^i + \pi_r U_{t+1}^r] \\ \text{Recovered:} & \quad U_t^r = u(c_t^r, n_t^r) + \beta U_{t+1}^r \end{aligned} \tag{1}$$

Here, τ_t represents the agents' probability of getting infected given their own and the infecteds' consumption and working activities, π_r is the probability of recovering, and π_d is the probability of death. We note

³In the pure SIR model, the median R_0 estimate in Liu et al. (2020) would result in an implausible 95% of the population contracting the disease.

here the strong assumption that agents know the ‘true’ infection probabilities. In Section 4.6, we illustrate the effects of perceived transmission probabilities on behavioral adjustment and agreement with lockdown policies.

3.1 Calibration

In the analysis below, we compare our calibration to the United States to one for Uganda as our example of a developing country. A first key difference between rich and poor countries is the lower IFR in the latter, mostly driven by their younger population and even after adjusting for health system differences (Ghisolfi et al., 2020; Walker et al., 2020). We use our predicted IFRs of 0.33% for Uganda and 0.79% for the United States, noting that ERT assumed an IFR of 0.5%. The most pessimistic scenario in Ghisolfi et al. (2020) accounts for the possibility of health systems becoming overburdened by assuming that none of the predicted severe cases can be saved in hospitals. We present below a robustness check of our main result using this adjusted IFR of 1.3%, but also note that recent evidence points in the opposite direction of potentially much lower IFRs in the context of developing countries.⁴

A second obvious difference between the two settings are incomes and hours worked. For the United States, we target a yearly income of \$58,000, earned during 28 hours of weekly labor, following ERT. For Uganda, we take the median main job monthly nominal wage for wage employees from the 2016/17 Household Survey (UBOS, 2018), converted to yearly dollar wages - \$535 - and set weekly hours worked to 50, following evidence for Tanzania and Ethiopia (Charmes, 2015). We set the subsistence level at \$200 per year in Uganda, the median nominal wage for female workers without formal education. In case of the United States, we stick to a formulation without a subsistence constraint as average incomes there are much further from such a level than Ugandan ones.⁵

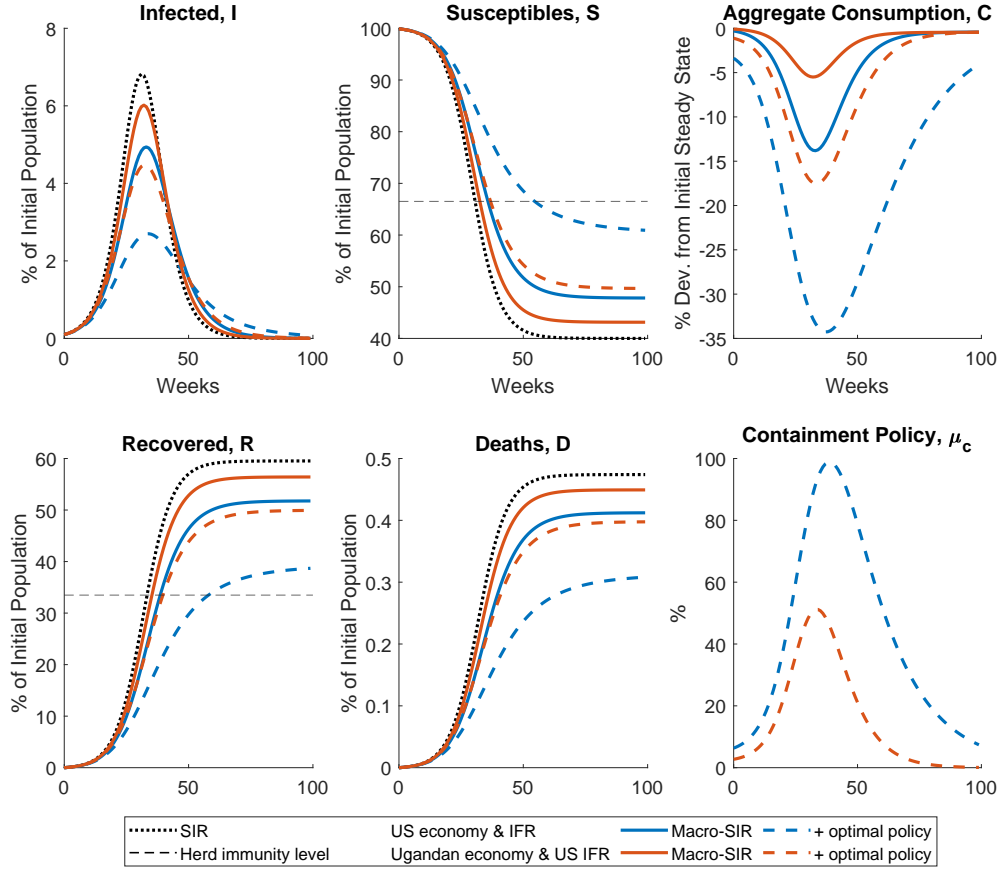
4 Results

We first present results from our calibration to the United States as a benchmark and explain the basic mechanisms of the model. We then take steps towards the low-income country calibration. We first change the income targets and subsistence level to the Ugandan values described above, holding constant the IFR.

⁴Seroprevalence studies from representative samples and reliable counts of fatalities would allow us to validate our estimates for developing countries. While we are not aware of such data yet, evidence from African countries at similar levels of development as Uganda supports our claim of generally lower IFRs in poorer countries by showing a high prevalence of COVID-19 antibodies in subgroups of the population (health personnel in Malawi: Chibwana et al. (2020); blood donors in Kenya: Uyoga et al. (2020)). In Kenya, the authors argue that if the disease was similarly deadly there as in other countries, official death counts would have to be implausibly inaccurate, implying that the IFR in Kenya is likely much lower than in other contexts.

⁵All other parameters are from ERT, including an annualized discount rate of 0.96, a 20% loss in productivity if infected, a recovery duration π_r of 18 days, one time unit representing a week, a model horizon of 250 weeks and an initial infection share of 1 in 1,000 people.

Figure 2: Calibrations to United States and Ugandan economy with United States IFR



The figure shows the time path of epidemiological (left and middle panels) and economic components (right panels) of the ERT model calibrated to the United States (blue) and Uganda (orange), holding constant the IFR. Dotted black line reports results from the basic SIR model. Solid lines depict agents' voluntary adjustments. Dashed lines show results from social planner problem maximizing overall utility.

The second step explores the role of a lower probability of deaths and subsequent lower aggregate disutility of infection/exposure. We then show how our results depend on the VSL we choose, before highlighting the role of beliefs about transmission probabilities in agreement with lockdown policies.

4.1 Optimal policy in the United States

The dotted black line in Figure 2 represents the course of the epidemic without any behavioral adjustment or government intervention - the basic SIR model. Infections grow until a large enough share of the population has acquired immunity, such that the infected become less likely to interact with the remaining susceptibles. At the end of the epidemic, 60% have ever been infected (targeted by the parameterization), and 0.47% have died (60% infected * 0.79% IFR = 0.47%). In our model, the epidemic ends when a sufficient share of the population has acquired immunity, the herd immunity level. However, the rapid increase in infections in the

pure SIR-model leads to an overshoot of infections (Moll, 2020). Hence, the final epidemic size as measured by the share of people that ever got infected lies above the minimal level necessary for herd immunity, indicated by the dashed horizontal line.

The solid blue line presents the SIR-model augmented with rationally adjusting agents. Focusing on the top-right panel, reductions in aggregate consumption amount to up to 13% and follow the infection rate. This is for two reasons: firstly, infected individuals are assumed to be 20% less productive and thus have less income to consume. With a peak infection rate of 5%, this amounts to a reduction of 1%. The larger share of the reduction comes from agents' voluntary adjustments of hours worked and consumption, in order to reduce the risk of infection. These adjustments slow the epidemic and reduce its peak infection rate, leading to a 13% reduction in death rates compared to the basic SIR model.

Besides this voluntary adjustment, the social planner can increase overall utility by internalizing the individual agents' contribution to the overall epidemic. This optimal policy, the dashed blue line in the bottom right panel, amounts to a tax on consumption equal to up to 100% at the peak of the epidemic or 56% over the first year, leading to an additional reduction in consumption of up to 35%. This slows the epidemic further, reducing deaths by an additional 21% and thus almost closing the gap between final epidemic size and the herd immunity level.


Overall, the United States calibration suggests that the policymaker can be effective at reducing deaths by reducing economic activity, and that this is socially optimal. Notably, the calculation implies that overall societal utility of the living between the voluntary adjustment and the optimal policy scenario is virtually unchanged, since agents are willing to reduce their working hours proportionately to the consumption reduction. This observation motivates our introduction of a subsistence constraint into the utility function, thus making the income elasticity of consumption dependent on baseline consumption.

4.2 The role of income levels in determining the optimal policy

Next, we assess how the economic outcomes of the model differ when income levels are lower (orange lines in Figure 2). Agents are now much poorer, less productive and face a non-zero subsistence constraint, but have the same, relatively high, probability of dying once infected as in the United States calibration.⁶ Focusing first on aggregate consumption adjustments without any containment policies, it is striking that the adjustments are less than half as strong than in the United States economy. The adaptation still reduces peak infection and death rates substantially, though less so than in the United States (8.8% reduction in death rate vs. 13%). We make a similar observation for the optimal policy, which now peaks at 41%: It

⁶The assumed fixed share of transmission through working or consuming explains the overlapping of the SIR model for either economy.

Figure 3: Calibration to Ugandan economy with Ugandan IFR



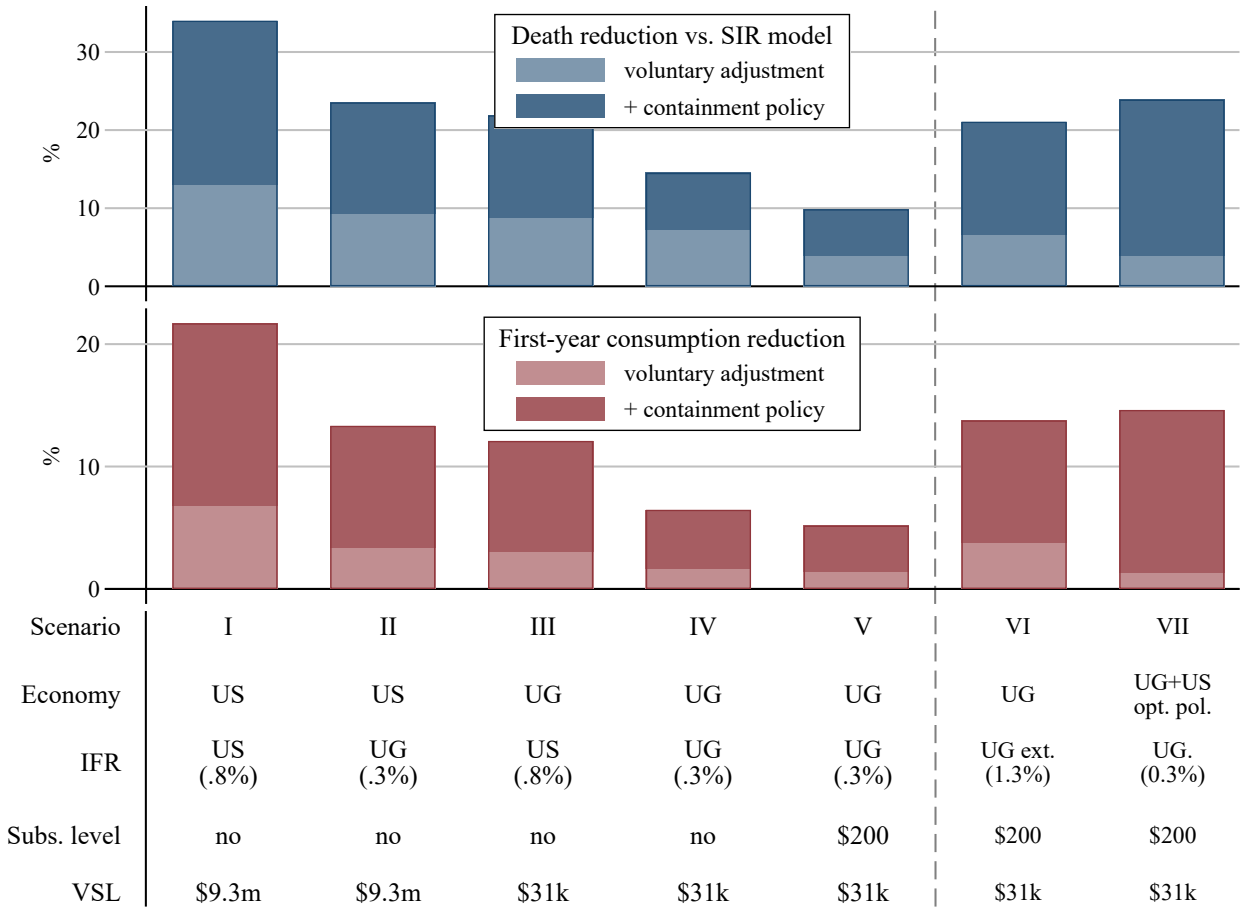
The figure shows the time path of epidemiological (left and middle panels) and economic components (right panels) of the model calibrated to Uganda with the United States IFR (orange) and Ugandan IFR (green). Dotted lines report results from the basic SIR model. Solid lines include agents' voluntary adjustments. Dashed lines show results from social planner problem maximizing overall utility.

reduces deaths, but much less so than in the United States (additional 13% reduction vs. 21%). Both agents and the social planner are trading off mortality risk and utility losses from consumption reductions. When consumption is low already it becomes more costly to reduce it, reductions lead to relatively large losses in utility, and it becomes less optimal to avert deaths. In other words, the externality from one agent getting infected, and not internalizing that she may infect others, becomes smaller relative to baseline utility the poorer the agent is.

4.3 The role of mortality risk in determining the optimal policy

The second step of the calibration keeps the Ugandan economy structure, and adds the estimated IFR for Uganda (green lines in Figure 3). Faced with a lower risk of death, agents reduce consumption only marginally, reducing the death rate by just 7.2%. Optimal policy now peaks at 21% and reduces deaths by

Figure 4: Optimal policy and death reductions along calibration steps from United States to Uganda



The figure shows summary statistics of voluntary adjustment and optimal policy implications under various calibration steps. The top panel in blue shows percentage reduction in deaths versus the underlying SIR model. The bottom panel in red shows decreases in consumption over the first year of the epidemic, relative to pre-epidemic levels. Light colors indicate effects of voluntary adjustments by agents trading off consumption and infection risk. Dark colors indicate additional effects of utility-maximizing containment policy. Scenarios differ by the economy and demography which they are calibrated to (Uganda and United States) and whether the utility function includes a subsistence level. Right-most column adds scenarios including the upper bound for the IFR from Ghisolfi et al. (2020) and applying United States optimal policy to our Ugandan calibration.

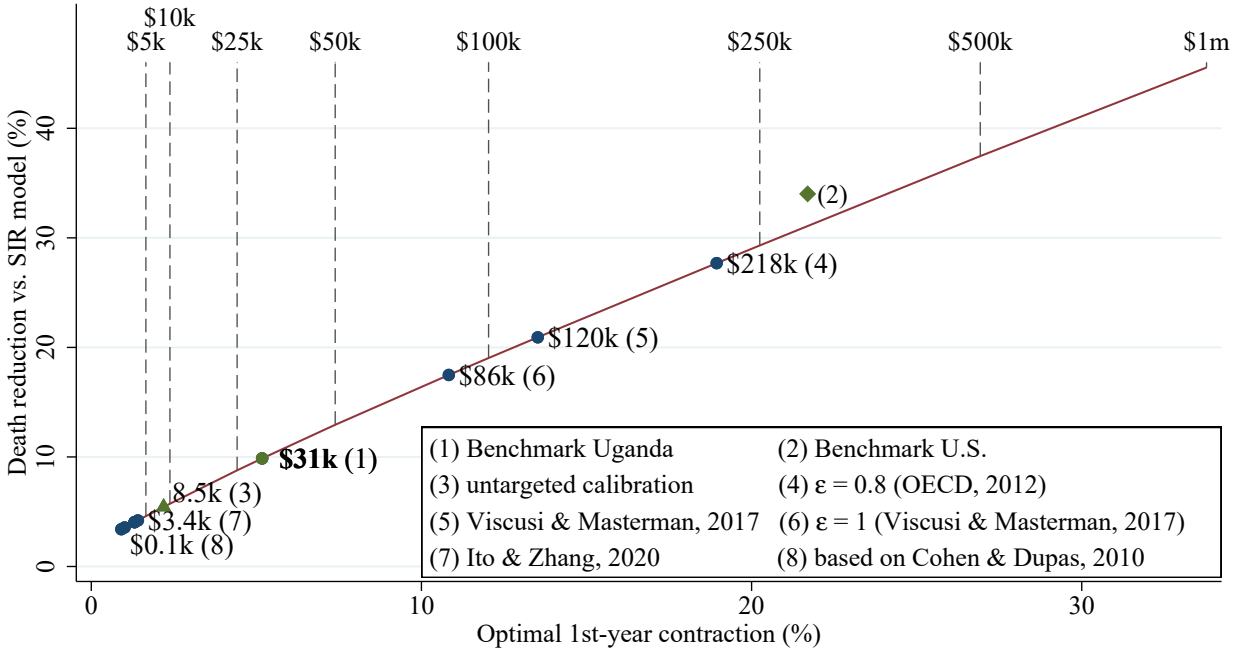
only an additional 7%, with final epidemic size far above the minimal herd immunity level.

4.4 Summary of calibration exercise

Our exercise highlights that even with equal death risk, the social planner should choose less stringent containment measures in a poorer economy, where reduction in consumption is relatively more costly in utility terms. Figure 4 shows the decrease in death reductions and consumption containment under optimal policies as we move from the United States towards to the Ugandan calibration.

As we change only the IFR between the first and second scenario, equivalent to moving from the 52nd to the 4th percentile of country-IFRs in Ghisolfi et al. (2020), optimal policy reduces deaths by 10 percentage

Figure 5: Optimal policy and death reductions for a range of assumed values of a statistical life



The red line traces out optimal policy effects on the economy (x-axis) and the epidemic (y-axis) for Uganda, depending on the adopted VSL. The assumed VSL rises the further from the origin. The blue dots identify other possible values as referenced below, the green diamond indicates death reduction and recession from benchmark U.S. calibration. The green triangle indicates the VSL from following ERT and setting $\bar{u} = 0$.

points less relative to the SIR model. Changing the economic structure to Uganda holding the IFR constant as in Figure 2 lowers the death reduction by a similar amount. Furthermore, our main qualitative result also holds when we assume a much higher death rate in Uganda, equivalent to all predicted severe cases succumbing to the disease (Scenario VI).

Finally, we can ask about the efficiency of a given policy in preventing deaths per unit of GDP reduction by transferring the optimal policy from the United States (Scenario I) to the Ugandan setup (Scenario VII). The voluntary reduction now constitutes a much smaller part of the total reaction, as a relatively strict lockdown is applied to agents not willing to incur large losses. The containment policy reduces deaths per 100,000 from 200 in the SIR model to 152, while reducing GDP by 14%. This yields a reduction in deaths per percentage point of GDP loss of 3.4, in contrast to the United States scenario where the same policy avoids 7.6 deaths per point of GDP loss.

4.5 The value of a statistical life as a critical parameter

As discussed in Section 2, our benchmark value for the cost of death falls into a broad range of possible values. Values that are high relative to agents' income approximate a view that lives should be saved even

at very high costs. At the other end of the scale, studies using revealed preference methods on peoples' own WTP to reduce and avoid fatality risks have found VSLs that are low relative to agents' income.

This wide range of values has implications for optimal policy in our framework, as presented in Figure 5. However, even when applying the VSL estimate based on the OECD (2012) study, optimal policy in Uganda would still be less strict as in the United States. If basing containment measures on the revealed preferences from the development economics literature, only very mild containment would be optimal. Our reference value of \$31,000 lies in between these estimates and optimal policy is thus moderately strict.

4.6 Beliefs about transmission risk and the acceptance of lockdown policies

Our finding on more muted optimal policies stands in contrast to the strict lockdown policies imposed in many developing countries, among them Uganda. Given the large difference, it may be surprising that recent surveys in Senegal and Pakistan have found broad agreement with the measures mandated by governments, despite households reporting substantial reductions in income (Moscoviz and Le Nestour, 2020; Brac, 2020).⁷

A possible explanation for these findings lies in that people may or would have been reducing their economic activity even without stricter government measures, as suggested by the mobility data from the United States and Sweden presented in Farboodi et al. (2020). Also in the case of Uganda, Mahmud and Riley (2020) report a large increase in protective behavior. In particular, this may be the case if agents overestimate the risk of getting infected and/or of dying from an infection. Further restrictions would then either not be controversial – if their effect does not exceed the voluntary reductions–, or even welcome – if there is a belief that others are not reducing their activity enough.⁸ With a small extension, our model lends itself to an analysis of the necessary overestimation. We here focus on the infection risk, but could perform a similar exercise for an overestimation of the death risk.

The 'true' infection risks from economic activity - π_{s1}, π_{s2} - are, in reality, unlikely to be known by agents. At least during early stages of the epidemic, beliefs likely often overstate the true (unknown) parameters in both developing and developed countries. Within the model, a 'fear parameter' ρ by which agents overestimate the true infection parameters from consumption or work captures this reasoning, transforming Equation 1 into

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

$$\text{with } \tau_t = \rho\pi_{s1}c_t^s (I_t C_t^I) + \rho\pi_{s2}n_t^s (I_t N_t^I) + \pi_{s3}I_t$$

⁷70% of respondents of a nationally representative survey in Senegal in mid-April supported a two-week lockdown when there had already been a nightly curfew and public spaces closed. This was similar among people who had already lost income, and higher among those more worried about the epidemic in general. In Bangladesh, respondents reported labor income decreases of up to 75%, while also generally supporting further restrictive measures.

⁸Mahmud and Riley (2020) report that only 14% of households in their rural Ugandan sample expect a member to contract the virus - absent representative testing in Uganda, we cannot say whether this is an over- or underestimate.

This exercise suggests that if agents overestimate the infection risk from consuming and working by factor six, they voluntarily reduce consumption by 5%, equalling the optimal adjustment. In turn, this illustrates that agreement with strict - according to our model overly strict - lockdown policies in developing countries can be partly explained through an overestimation of the risks. Governments may hence find it harder to restrict economic activity once people form more accurate or even too optimistic beliefs about the medical risks.

Another, related explanation for strict initial lockdowns could be that governments used containment policies as signalling tools to individuals, foreseeing that compliance would be incomplete if the risks were objectively known and enforcement capacity limited. In our framework, we can think of this as the policy variable not only containing behavior directly through disincentivizing consumption, but also through shaping, at least initially, beliefs about risks. The extent of such effects would be an interesting avenue for future research.

5 Conclusion

We integrate behavioral responses in a standard SIR epidemiological model by building on Eichenbaum et al. (2020) and extend their model to take account of subsistence constraints as well as demography-dependent fatality rates. We use the model to compare optimal policy responses to the current pandemic for two countries, the United States and Uganda, as stylized examples of countries varying by poverty and demography. We calibrate the model with country-specific distributions of age, comorbidities, and income, extend the framework to allow for a ‘fear parameter’ that may induce agents to reduce economic activity more than socially optimal given the diseases’ true risks, and highlight the dependence of our and related studies on the chosen value of the cost of death, which may fall into a broad range.

The differences in *optimal* policy responses we find between the United States and Uganda stand in contrast to the relatively similar *actual* initial policy responses across the world. Our results suggest that lockdown policies in Uganda and other countries with similar income levels and demographics may be too restrictive compared to optimal policies. This begs the question whether governments are optimizing as social planners and may have gotten it wrong, or whether their actions can be rationalized with considerations outside our model. One simple explanation would be that in the face of huge uncertainty, governments have adopted approaches from other countries which had already gained more experience with the epidemic, without adapting them to local conditions. As an alternative, we explored the role of possible overestimation of the fatality risk among the population, leading to demand for stronger measures. Both explanations are in line with the observation that restrictions have been scaled down the fastest in developing countries (see

Figure 3). Another possible explanation is that governments in poorer countries placed a high value on protecting lives of their citizens, possibly higher relative to their GDP than a standard value in the United States.

Finally, we reiterate that a central assumption to our analysis is that the societal disutility of death equals the foregone utility of living. While common in macroeconomic research, there may well be different societal preferences underlying the choices taken by governments during this time. As Ghana's President Nana Akufo-Addo said in March explaining lockdown measures, "We know how to bring the economy back to life. What we do not know is how to bring people back to life. We will, therefore, protect people's lives, then their livelihoods." An interesting topic for future research would be to consider potential heterogeneities in country preferences that could rationalize the actual responses we observe.

References

- Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D., 2020. Optimal Targeted Lockdowns in a Multi-Group SIR Model. NBER Working Paper 27102. doi:10.3386/w27102.
- Adler, M.D., 2020. What should we spend to save lives in a pandemic? A critique of the value of statistical life. CEPR Covid Economics 33.
- Alon, T., Kim, M., Lagakos, D., VanVuren, M., 2020. How Should Policy Responses to the COVID-19 Pandemic Differ in the Developing World? NBER Working Paper 27273. doi:10.3386/w27273.
- Alvarez, F.E., Argente, D., Lippi, F., 2020. A Simple Planning Problem for COVID-19 Lockdown. NBER Working Paper 26981. doi:10.3386/w26981.
- Avery, C., Bossert, W., Clark, A., Ellison, G., Ellison, S.F., 2020. Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists. NBER Working Paper 27007. doi:10.3386/w27007.
- Barnett-Howell, Z., Mobarak, A.M., 2020. The Benefits and Costs of Social Distancing in Rich and Poor Countries. Working Paper arXiv:2004.04867.
- Berry, J., Fischer, G., Guiteras, R., 2020. Eliciting and utilizing willingness to pay: Evidence from field trials in northern Ghana. *Journal of Political Economy* 128, 1436–1473. doi:10.1086/705374.
- Brac, 2020. Rapid Perception Survey On COVID19 Awareness and Economic Impact. Working Paper .
- Charmes, J., 2015. Time use across the world: Findings of a world compilation of time use surveys. UNDP Human Development Report Office .
- Chetty, R., Friedman, J.N., Hendren, N., Stepner, M., Opportunity Insights Team, 2020. How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data. NBER Working Paper doi:10.3386/w27431.
- Chibwana, M., Jere, K., Kamng'ona, R., Mandolo, J., Katunga-Phiri, V., Tembo, D., Mitole, N., Musasa, S., Sichone, S., Lakudzala, A., Sibale, L., Matambo, P., Kadwala, I., Byrne, R., Mbewe, A., Henrion, M., Morton, B., Phiri, C., Mallewa, J., Mwandumba, H., Adams, E., Gordon, S., Jambo, K., 2020. High SARS-CoV-2 seroprevalence in health care workers but relatively low numbers of deaths in urban Malawi [version 1; peer review: awaiting peer review]. Wellcome Open Research doi:10.12688/wellcomeopenres.16188.1.

- Chudik, A., Pesaran, M.H., Rebucci, A., 2020. Voluntary and Mandatory Social Distancing: Evidence on COVID-19 Exposure Rates from Chinese Provinces and Selected Countries. NBER Working Paper 27039. doi:10.3386/w27039.
- Cohen, J., Dupas, P., 2010. Free Distribution or Cost-Sharing? Evidence from a randomized Malaria Prevention Experiment. *Quarterly Journal of Economics* 75. doi:10.1093/qje/qjs020.Advance.
- Desai, S., Pramanik, S., 2020. Virus fear and logistical challenges curtailing economic activity in Delhi-NCR, new survey shows. *livemint* URL: bit.ly/livemint-dcvts.
- Eichenbaum, M., Rebelo, S., Trabandt, M., 2020. The Macroeconomics of Pandemics. CEPR Discussion Paper 14520. doi:10.1017/CBO9781107415324.004.
- Farboodi, M., Jarosch, G., Shimer, R., 2020. Internal and External Effects of Social Distancing in a Pandemic. CEPR Covid Economics 9.
- Garibaldi, P., Moen, E., Pissarides, C., 2020. Modelling contacts and transitions in the SIR epidemics model. CEPR Covid Economics 5.
- Ghisolfi, S., Almås, I., Sandefur, J., von Carnap, T., Bold, T., 2020. Calculating regional differences in COVID-19 fatality rates. *BMJ Global Health* 5. doi:10.1136/bmjgh-2020-003094.
- Glover, A., Heathcote, J., Krueger, D., Ríos-Rull, J.V., 2020. Health versus Wealth: On the Distributional Effects of Controlling a Pandemic. CEPR Covid Economics 6.
- Hale, T., Webster, S., Petherick, A., Phillips, T., Kira, B., 2020. Oxford COVID-19 Government Response Tracker, Blavatnik School of Government.
- Hausmann, R., Schetter, U., 2020. Horrible Trade-offs in a Pandemic: Lockdowns, Transfers, Fiscal Space, and Compliance. Working Paper doi:10.5874/jfsr.19.225.
- IGM Forum, 2020. Policy for the COVID-19 Crisis. URL: <http://www.igmchicago.org/surveys/policy-for-the-covid-19-crisis/>.
- Ito, K., Zhang, S., 2020. Willingness to pay for clean air: Evidence from air purifier markets in China. *Journal of Political Economy* 128, 1627–1672. doi:10.1086/705554.
- Kremer, M., Leino, J., Miguel, E., Zwane, A.P., 2011. Spring cleaning: Rural water impacts, valuation, and property rights institutions. *Quarterly Journal of Economics* 126, 145–205. doi:10.1093/qje/qjq010.

- Krueger, D., Uhlig, H., Xie, T., 2020. Macroeconomic dynamics and reallocation in an epidemic. CEPR Covid Economics 5.
- Liu, Y., Gayle, A.A., Wilder-Smith, A., Rocklöv, J., 2020. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine* 27, 1–4. doi:10.1093/jtm/taaa021.
- Mahmud, M., Riley, E., 2020. Household response to an extreme shock : Evidence on the immediate impact of the Covid-19 lockdown on economic outcomes and well-being in rural Uganda. Working Paper .
- Moll, B., 2020. Lockdowns in SIR Models. URL: https://benjaminmoll.com/wp-content/uploads/2020/05/SIR_notes.pdf
- Moscoviz, L., Le Nestour, A., 2020. Five Findings from a New Phone Survey in Senegal URL: <https://www.cgdev.org/blog/five-findings-new-phone-survey-senegal>.
- Nordling, L., 2020. The pandemic appears to have spared Africa so far. Scientists are struggling to explain why (August 7th, 2020). *Science Magazine* .
- OECD, 2012. Mortality Risk Valuation in Environment, Health and Transport Policies. OECD Publishing.
- Ravallion, M., 2020. Pandemic Policies in Poor Places. CGD Note URL: <https://www.cgdev.org/sites/default/files/pandemics-in-poor-places.pdf>.
- Ray, D., Subramanian, S., 2020. India's Lockdown : An Interim Report. NBER Working Paper 27282.
- Ray, D., Subramanian, S., Vandewalle, L., 2020. India's Lockdown. CEPR Policy Insight .
- Robinson, L.A., Hammitt, J.K., O'Keeffe, L., 2019. Valuing Mortality Risk Reductions in Global Benefit-Cost Analysis. *Journal of Benefit-Cost Analysis* 10, 15–50. doi:10.1017/bca.2018.26.
- Toxvaerd, F., 2020. Equilibrium social distancing. CEPR Covid Economics 15.
- UBOS, 2018. Uganda National Household Survey Report 2016 / 2017 URL: <http://www.ubos.org>.
- Uyoga, S., Adetifa, I.M., Karanja, H.K., Nyagwange, J., Tuju, J., Wanjiku, P., Aman, R., Mwangangi, M., Amoth, P., Kasera, K., Ng'ang'a, W., Rombo, C., Yegon, C.K., Kithi, K., Odhiambo, E., Rotich, T., Orgut, I., Kihara, S., Otiende, M., Bottomley, C., Mupe, Z.N., Kagucia, E.W., Gallagher, K., Etyang, A., Voller, S., Gitonga, J., Mugo, D., Agoti, C.N., Otieno, E., Ndwiga, L., Lambe, T., Wright, D., Barasa, E., Tsofa, B., Bejon, P., Ochola-Oyier, L.I., Agweyu, A., Scott, J.A.G., Warimwe, G.M., 2020. Seroprevalence of anti-SARS-CoV-2 IgG antibodies in Kenyan blood donors. medRxiv doi:10.1101/2020.07.27.20162693.
- Viscusi, W., Masterman, C.J., 2017. Income Elasticities and Global Values of a Statistical Life. *Journal of Benefit-Cost Analysis* 8, 226–250. doi:10.1017/bca.2017.12.

Walker, P.G., Whittaker, C., Watson, O.J., Baguelin, M., Winskill, P., Hamlet, A., Djafaara, B.A., Cucunubá, Z., Mesa, D.O., Green, W., Thompson, H., Nayagam, S., Ainslie, K.E., Bhatia, S., Bhatt, S., Boonyasiri, A., Boyd, O., Brazeau, N.F., Cattarino, L., Cuomo-Dannenburg, G., Dighe, A., Donnelly, C.A., Dorigatti, I., Van Elsland, S.L., FitzJohn, R., Fu, H., Gaythorpe, K.A., Geidelberg, L., Grassly, N., Haw, D., Hayes, S., Hinsley, W., Imai, N., Jorgensen, D., Knock, E., Laydon, D., Mishra, S., Nedjati-Gilani, G., Okell, L.C., Unwin, H.J., Verity, R., Vollmer, M., Walters, C.E., Wang, H., Wang, Y., Xi, X., Lalloo, D.G., Ferguson, N.M., Ghani, A.C., 2020. The impact of COVID-19 and strategies for mitigation and suppression in low- and middle-income countries. *Science* 369, 413–422. doi:10.1126/science.abc0035.