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

Tillmann von Carnap[a,b] Ingvild Almås[a,c] Tessa Bold[a]

Selene Ghisolfi[a,d] Justin Sandefur[e]

Abstract

How do optimal policies to control the spread of SARS-CoV-2 vary across countries? In an influential recent paper, Eichenbaum, Rebelo, and Trabandt (2020) incorporate economic behavior into a standard epidemiological model calibrated to the United States, finding that spontaneous social distancing will fall short of the social optimum without policy intervention. In this paper, we apply and extend their model to explore how optimal policy varies across contexts depending on demography, comorbidities, and health system strength – which affect the infection fatality rate – as well as poverty – which affects agents' willingness to forego current consumption to reduce disease risk. Calibrating the model to Uganda, we calculate an optimal path for a containment policy equivalent to a 4% consumption tax over one year (compared to a 40% tax in the U.S.), which reduces predicted mortality by roughly 5.5% (compared to 28.2% in the U.S.). Notably, the normative predictions of the model constitute poor positive predictions. Actual containment policies in Uganda and many other developing countries with high poverty and favorable demographics have been relatively severe, and have been met with broad public approval despite high economic costs. Within the model, widespread overestimation of the risk of contracting and dying from the disease provides one possible explanation for this pattern.

[a]Institute for International Economic Studies, Stockholm University
[b]Corresponding author: tillmann.voncarnap@iies.su.se
[c]Norwegian School of Economics
[d]LEAP, Bocconi University
[e]Center for Global Development

1 Introduction

The ongoing COVID-19 pandemic has led governments around the world to impose unprecedented restrictions on economic activity, with surprising uniformity across countries at all income levels (Figure 1). While economists in rich countries have largely spoken out in support of strict containment policies on the grounds that controlling the pandemic is a prerequisite for economic recovery (IGM Forum, 2020; Mahoney, 2020; Summers, 2020), development economists have expressed reservations about the paths taken by governments in poorer countries (Ray and Subramanian, 2020; Ray et al., 2020; Barnett-Howell and Mobarak, 2020; Ravallion, 2020). We explore two reasons why a unified economic framework might produce such different conclusions: demography and poverty.

We build on the recent model developed by Eichenbaum, Rebelo, and Trabandt (2020, henceforth ERT) to compare welfare-optimizing policies across contexts that differ in terms of these two factors. We replicate the authors' analysis of optimal policy in the United States, and compare it to results calibrated to Uganda. We focus on the latter as an example of a developing country highlighting the two salient features we seek to explore. First, Uganda's median age of 16 is less than half than

Figure 1: OxCGRT Stringency Index in 2020 - selected countries and average per income group

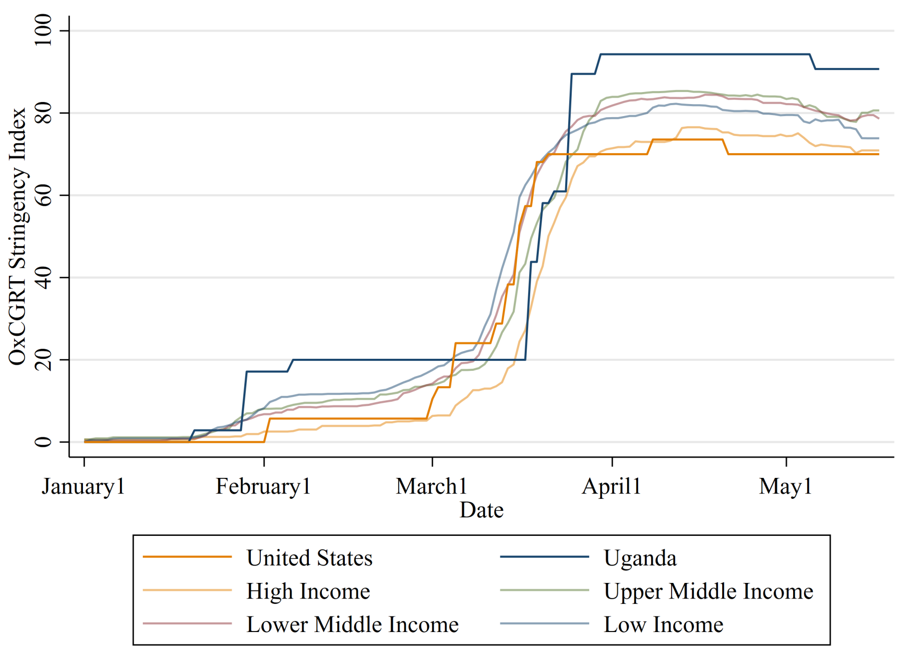

The figure shows the average level of restrictions over time in the four income groups defined by the World Bank and our study countries as measured by the Oxford Coronavirus Government Response Stringency Index. The index 'systematically collects information on several different common policy responses that governments have taken to respond to the pandemic on 17 indicators such as school closures and travel restriction'.

Figure 2: IFR estimates for US and Uganda, based on country-specific age and comorbidity distributions

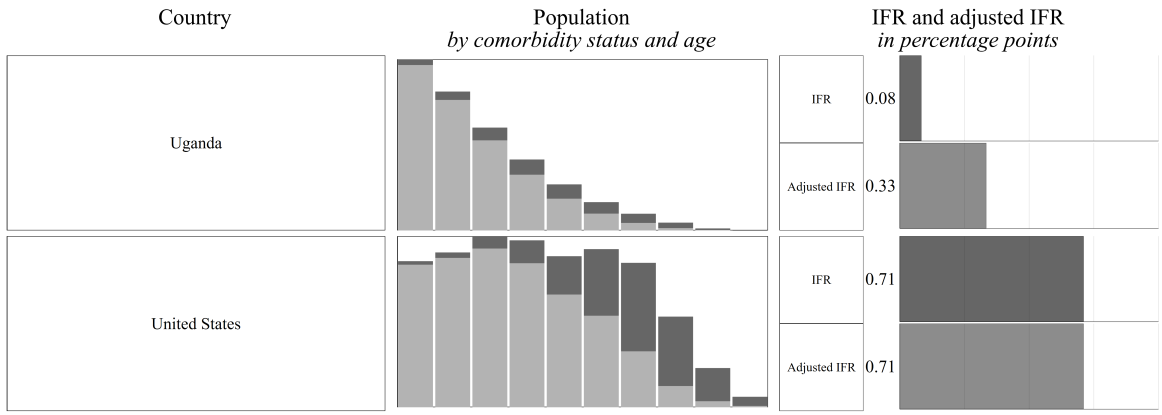

The figure shows the IFR estimates for US and Uganda, based on country-specific age and comorbidity distributions. The middle column reports population by 10-year age groups and by number of comorbidities (light grey - 0 comorbidities; dark grey - any comorbidity), the height of the bars is proportional to the number of people in the most populous age group. Dark grey bars in the right column report IFRs, calculated as an average of the IFRs by age and comorbidity weighed by the proportion of the population in each age and comorbidity group. Light grey bars show the IFR taking into account local health system capacity. See Ghisolfi et al. (2020) for further reference.

that of the US, implying the share of infected who die from the disease may be considerably lower.[1] In the ERT framework, this risk differential affects agents' labor supply and consumption decisions, and in turn optimal policy. Second, the country's GDP per capita at $710 is a fraction of the US's at approximately $54,000.[2] At lower incomes foregone consumption due to pandemic control (voluntary or policy-induced) implies a larger welfare loss in the ERT framework, affecting agents' and policymakers' optimal choices. Finally, despite these differences, Uganda's response to the virus has been relatively early and strict (see Figure 1), and available figures suggest successful containment: by early June, there had been 500 confirmed cases and no reported deaths[3].

A key statistic for our analysis summarizing the above demographic differences is the infection fatality ratio (IFR), the share of people dying from an infection. In companion work, we show that a population's age and comorbidity structure predicts wide variation in this statistic across contexts (Ghisolfi et al., 2020). Even after adjusting for the lower capacity of health systems in Uganda than in the United States, we find that an infected person in Uganda is less than half as likely to die from COVID-19 compared to the US (Figure 2). We describe the construction of these numbers in Section 2.1.

In recent weeks, economists have put forward various ways to incorporate economic decision-making

[1] For comparison, the average is 20 years in low-income countries and 40 in high-income countries. Source: https://apps.who.int/gho/data/view.main.POP2030

[2] The average among low-income countries according to the World Bank's World Development Indicators is $750, versus $43,000 among high-income countries.

[3] Source: ourworldindata.org, accessed June 4th, 2020

into standard epidemiological models, including the Susceptible-Infected-Recovered (SIR) framework. Studies have highlighted that voluntary adjustments of economic activity by agents facing contagion risk may go some way in containing the spread of the epidemic (Toxvaerd, 2020; Garibaldi et al., 2020; Chudik et al., 2020), and that, from a social welfare perspective, further government action can be 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 of these recent papers have focused on developing countries and the efficacy of the aforementioned mechanisms in their contexts (Alon et al. (2020) being a notable exception we discuss below). Our calibration exercise to a developing country, Uganda, focuses on inelastic consumption adjustments and implicit valuations of the utility of the living against the number of deaths. We further explore a simple extension to this model which can help explain the seemingly paradoxical observation from recent surveys in low-income countries, where respondents state high rates of agreement with lockdowns imposed by governments, while at the same time experiencing large income losses.

Our primary reference, the ERT model, posits a continuum of representative agents who value consumption and leisure. The agents expose themselves to infection risk when working and consuming, and, realizing this danger, reduce work hours and consumption as the risk of infection rises. However, they do not weigh the impact of their own labor and consumption decisions on the pandemic's spread, leading to an externality the social planner seeks to internalize. This is modeled via a discouragement on consumption, termed the 'containment rate' by the authors. The optimal policy here is one that maximizes the present value of societal utility. Components of this are i) the utility the agents derive from consumption, labor income, and leisure, ii) the disutility of foregone consumption from lower productivity when infected, and iii) the disutility of dying. The social planner thus strikes a balance between lost utility from containment policies, and lost utility from infection and death. In line with other studies, the authors show that the time path of 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. Although never restricting economic activity totally, their calibration yields a strong and sustained discouragement of consumption, equivalent to a 40% consumption tax over the first year of the epidemic.

Inherent in their analysis and most other recent modeling efforts of optimal lockdown policy is a trade-off between lives saved and foregone consumption. In the ERT framework, a central assumption, which we take on board here, 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, in all societies around the world, decision makers are making trade-offs that either 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 such as roads. While the qualitative results of our study do not depend on the precise value put on a life by either the agent or the social planner, our quantitative statements do. In future work, we plan to explore the implicit valuations put on lives by policymakers, but for now we adhere to the ERT assumption, which is standard in macroeconomic models.

We find that the difference in terms of mortality between a simple SIR model and one integrating economic decision making is much smaller in Uganda than in the case of the United States. While ERT find that the optimal containment policy in the United States reduces mortality by 28% relative to the pure SIR model, we find the corresponding figure for Uganda is only 5.5%. Both of our suggested factors contribute to this result: First, lower infection fatality ratios (IFR) 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 their labor supply and 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 the 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.

Given this discrepancy between observed and optimal policies, we highlight how a key assumption, that the agents are aware of the true transmission and death probabilities, may be adjusted to explain why acceptance of strict policies in many contexts appears surprisingly high: if agents perceive an exaggerated risk of contracting the disease and dying from it, their voluntary adjustment to protect themselves may coincide with the effects of the strict measures introduced by governments. Unless they updates their beliefs, agents may well agree with a painfully strict lockdown.

Our study complements independent research by Alon et al. (2020) who use a heterogenous agents model including various characteristics of developing country state and health system capacity to assess how different levels of governmental restrictions affect welfare in rich and poor countries. While their model assumes that workers in the informal sector cannot be shielded from the disease by a lockdown, the authors argue that government policy will be less effective in containing the epidemic in countries with larger informal sectors. Our approach highlights a similar effect, though grounded in the utility maximization of the representative agent: When faced with a relatively large decrease in consumption, a poorer agent will rationally reduce her exposure to the epidemic less, requiring government efforts that are stricter than in richer contexts 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 not only mechanically affects the overall death rate, but also individual-level optimization and adaptation behavior.

In the following section, we present in more detail the model by ERT and discuss our calibration procedure to the case of Uganda. In Section 3, we present the results before concluding in Section 4.

2 Model

Since the outbreak of the current pandemic, economists have quickly incorporated basic formulations of the SIR model into economic frameworks (Atkeson, 2020). In this, they have largely relied on an early formulation by Kermack and McKendrick (1927), 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. We first give a very short introduction to this class of models, and then discuss its incorporation into an economic framework.

The epidemic starts with some 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 a rate γ and are thereafter assumed to be immune. Depending on the relative size of β and γ, the epidemic dies 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. An extension also models the share of people that die from the epidemic as a simple share of those contracting the disease. We note that estimates of these parameters are still surrounded by significant uncertainty, and that the question of whether surviving an infection provides immunity is not settled (Avery et al., 2020).

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

$$T_t = \pi_{s1} \left(S_t C_t^s\right) \left(I_t C_t^i\right) + \pi_{s2} \left(S_t N_t^s\right) \left(I_t N_t^i\right) + \pi_{s3} S_t I_t \tag{1}$$

The first term on the right denotes new infections coming from interactions while consuming, or more broadly defined, spending money. This can include activities such as shopping, using services or leisure traveling. It is higher the more susceptible (S_t) and infected (I_t) people there are, and the more each of them consume $\left(C_t^s, C_t^i\right)$. Similarly, the second term describes infections from working, increasing in the number of hours worked by members of the susceptible and infected group $\left(N_t^s, N_t^i\right)$. 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. ERT assume that evidence from other epidemics also applies to the current one, in that 1/6 of infections take place at work or while consuming, and 2/3 from random interactions[4]. The authors furthermore target an immunity level of 60% in the population when the epidemic has run its course. Taking these targets, they calibrate the π_s using the steady state values of hours worked and consumption.

Turning to the economic side of the model, the economy consists of a continuum of representative agents whose behavior is modeled through a simple utility function, trading off consumption c and labor n.

$$u\left(c_t, n_t\right) = lnc_t - \frac{\theta}{2} n_t^2$$

Agents 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:

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

In choosing the parameter θ and the wage rate, ERT target two readily available statistics: the average weekly income and the number of hours worked. The parameters can then be calculated from the model's steady state.

Agents maximize their lifetime utility, the form of which differs by whether they are susceptible, infected, or recovered.

$$\begin{array}{lcll} \text{Susceptible:} & U_t^s & = & u\left(c_t^s, n_t^s\right) + \beta\left[\left(1-\tau_t\right) U_{t+1}^s + \tau_t U_{t+1}^i\right] \\ \text{with} & \tau_t & = & \pi_{s1} c_t^s \left(I_t C_t^I\right) + \pi_{s2} n_t^s \left(I_t N_t^I\right) + \pi_{s3} I_t \\ \text{Infected:} & U_t^i & = & u\left(c_t^i, n_t^i\right) + \beta\left[\left(1-\pi_r-\pi_d\right) U_{t+1}^i + \pi_r U_{t+1}^r\right] \\ \text{Recovered:} & U_t^r & = & u\left(c_t^r, n_t^r\right) + \beta U_{t+1}^r \end{array} \quad (2)$$

[4]The transmission probabilities in the ERT framework are, due to scarcity of reliable estimates, calibrated on a weak evidence base and on strong assumptions. In addition, it is unclear whether more or less infections in developing countries come through consumption and work: while crowded work places may well be more common in developing countries, work is also more likely to take place outdoors where transmission is lower. Absent better evidence, we have here copied the calibration exercise in ERT. We note that if transmissions from consumption and work are more important (and perceived so) in developing countries, the economic reactions modeled in this paper would be stronger, and vice versa. Collecting evidence on transmission through the various channels, and their perceptions would thus be an important avenue for future research.

Here, τ_t represents the agents' probability of getting infected given their own and the infected's consumption and working activities, π_r is the probability of recovering, and π_d is the probability of death. We note here the strong assumption that agents know the 'true' infection probabilities - it seems likely that different parts of the population in all countries may over- or underestimate this probability, potentially by a large factor. We present below an exercise illustrating the effects of varying perceived transmission probabilities on behavioral adjustment and agreement with lockdown policies. The formulation above furthermore expresses the assumption that the cost of death is the foregone utility of life. While common in macroeconomic research, it imposes a strict valuation of a life, on which optimal policy depends. As agents do not fully internalize their own contribution to the spread of the epidemic, the social planner can efficiently further discourage economic activity through the containment rate. In particular, ERT set the objective of the policy maker to maximize overall societal utility, including disutility of the dead.

2.1 Calibration

In our analysis below, we compare ERT's calibration to the United States to one for Uganda as our example of a developing country. As we show in companion research, a key difference between rich and poor countries is likely to lie in the lower share of infected succumbing to the disease in the latter, mostly driven by their younger and less vulnerable population (Ghisolfi et al., 2020). This difference remains even after adjusting for the capacity of health systems to deal with pulmonary diseases. We predict an IFR of 0.33% for Uganda and 0.71% for the United States. This figure takes into account the age and comorbidity distribution, as well as health system capacity of the countries. ERT assume an IFR of 0.5% for the US and for comparability reasons, we keep this rate for the US calibration. However, we can easily show qualitative robustness related to incorporating our estimated IFR of 0.71% for the United States. In fact, the differences in optimal policy we highlight are larger the greater the relative difference in IFRs between contexts.

An obvious difference between the two settings are incomes and hours worked. These are used to calibrate the disutility of labor parameter of the utility function and labor productivity. For the US, we again follow ERT, who target a yearly income of $58,000, earned during 28 hours of weekly labor. For Uganda, we take the median main job monthly nominal wage for wage employees from the 2016/17 Household Survey, converted to yearly dollar wages (UBOS, 2018), amounting to $535, and set weekly hours worked to 50, in line with evidence from Tanzania and Ethiopia presented in Charmes (2015). We take all other parameters from ERT. This includes an annualized discount rate of 0.96, and a recovery duration π_r of 18 days. Note that the model is set up with one time unit representing a week, so all parameters are converted accordingly. The model runs for 250 weeks, long enough to have the epidemic run its course. At the beginning of the epidemic, 1 in 1000 people are infected.

3 Results

In the following, we start by replicating Figure 2 and 3 from ERT to serve as a benchmark scenario and to explain the basic mechanisms of the model. We then proceed in steps towards a scenario calibrated to a low- income country setting. The first step is to change the income targets to the values described for Uganda in the calibration above, keeping the IFR at the US level. This comparison will highlight the role of income per se, holding constant the disutility of death relative to the utility of living. The second step, is to apply the IFR for Uganda and explore the role of a lower probability of deaths and subsequent lower aggregate disutility of infection/exposure. The third step is to introduce a subsistence constraint and by this deviate from ERT's formulation of the representative agents' utility function. Finally, we present a small extension highlighting the role of beliefs about transmission probabilities in agreement with lockdown policies.

3.1 Replicating ERT's calibration to the United States

Figure 3 replicates a combination of Figure 2 and 3 in ERT. The dotted black line represents the course of the epidemic without any behavioral adjustment or government intervention - the basic SIR model. As is common to these models, infections increase ever faster until a large enough share of the population has acquired immunity, such that the infected are less and less likely to interact with the remaining susceptibles. At the end of the epidemic, 60% of the population have ever been infected (targeted by the parameterization), and 0.3% have died (60% infected * 0.5% of the infected dead = 0.3% of the population).

The solid blue line presents the model estimates from augmenting the SIR-model with rationally adjusting agents. Focusing on the top-right panel, reductions in aggregate consumption amount to up to 8% 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 10.7% reduction in death rates.

Figure 3: Replication of ERT calibration to US economy

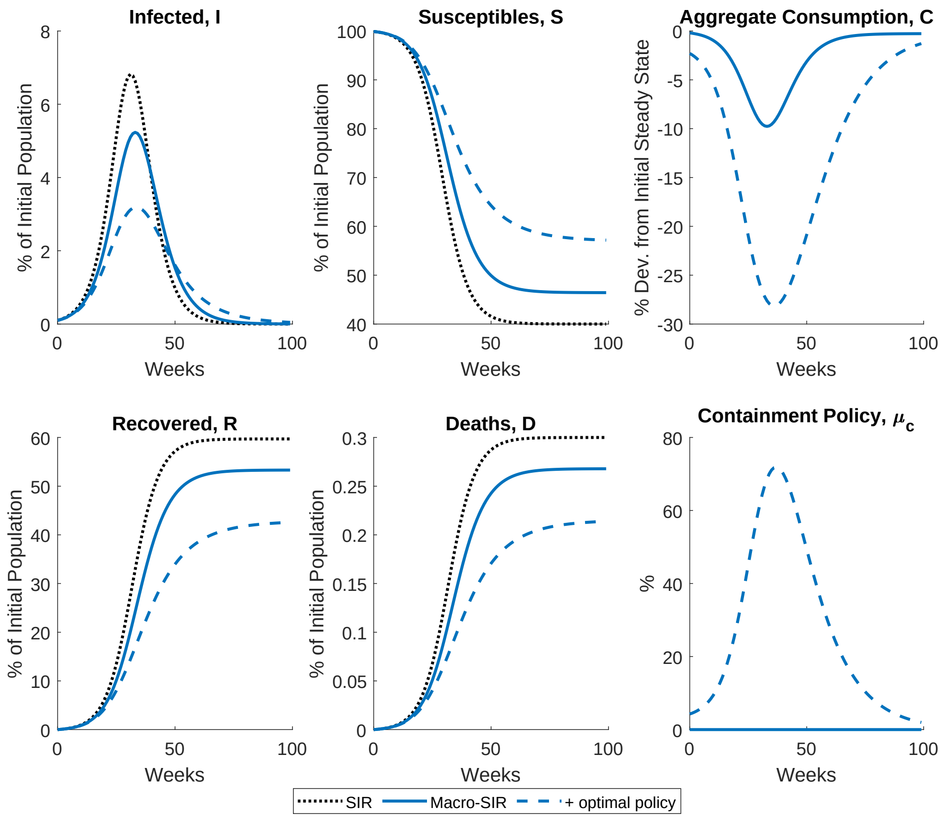

The figure shows the time path of epidemiological (left and middle panels) and economic components (right panels) of the ERT model calibrated to the United States. Dotted black line reports results from the basic SIR model. Solid line includes agents' voluntary adjustments. Dashed line shows results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

Besides this voluntary adjustment, the social planner can increase overall utility by internalizing the individual agents' contribution to the overall epidemic, which is otherwise not taken into account. This optimal policy, the dashed blue line in the bottom right panel amount to a tax on consumption equal to up to 70% at the peak of the epidemic or 40% over the first year, leading to an additional reduction in consumption of up to 28%. This slows the epidemic further, reducing deaths by an additional 17.5%.

Overall, the ERT calibration suggests that the policymaker can be effective at reducing deaths by reducing economic activity. Interestingly, their 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 exercise in Section 3.4 where we introduce a subsistence constraint into the utility function, thus making the income elasticity of consumption dependent on baseline consumption.

Figure 4: Calibration to Ugandan economy with US IFR

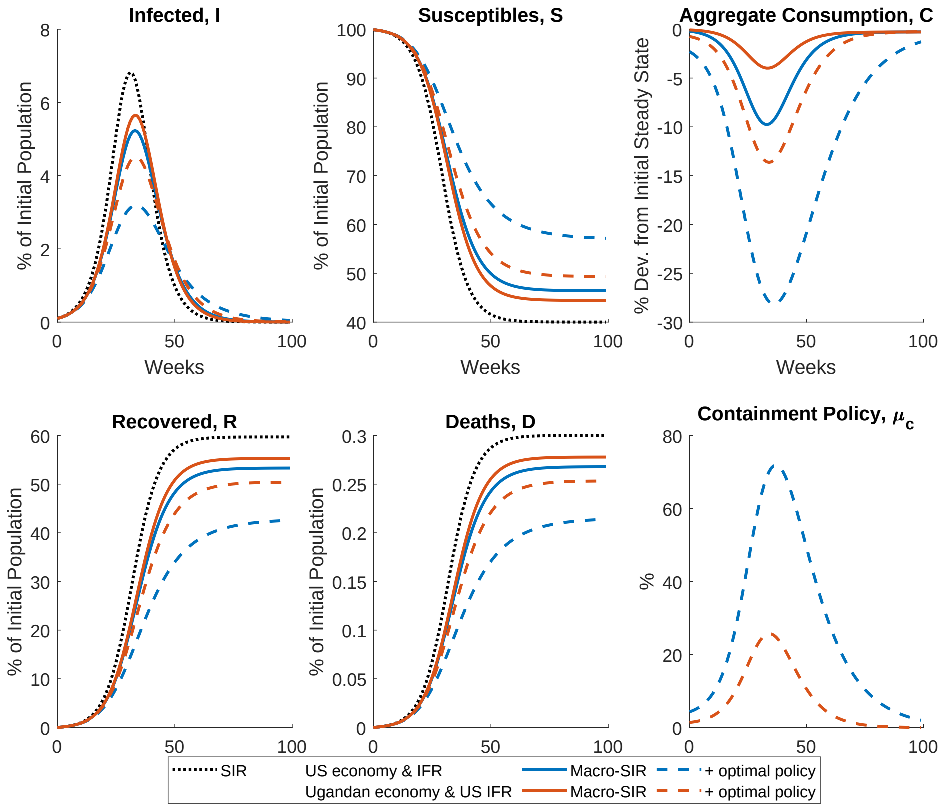

The figure shows the time path of epidemiological (left and middle panels) and economic components (right panels) of the ERT model calibrated to the United States (blue) and Uganda (orange), holding constant the IFR. Dotted black line reports results from the basic SIR model. Solid lines include agents'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

3.2 The role of income levels in determining the optimal policy

We now take the first step of our calibration to Uganda by changing the economic targets of the model. Agents are now much poorer (\$535 vs. \$58,000) and less productive (working 50hrs/week vs. 28), but have the same, relatively high, probability of dying once infected as in the US calibration. Figure 4 repeats the graphs from the US calibration in Figure 3 in blue, and adds the same set of graphs for the Ugandan economy in orange. 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 US economy. The adaptation still reduces peak infection and death rates substantially, though less so than in the US (7.4% reduction in death rate vs. 10.7%). We make a similar observation for the optimal policy, which now peaks at 26%: It reduces deaths, but much less so than in the US (additional 8.1% reduction vs. 17.5%). 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,

Figure 5: Calibration to Ugandan economy with Ugandan IFR

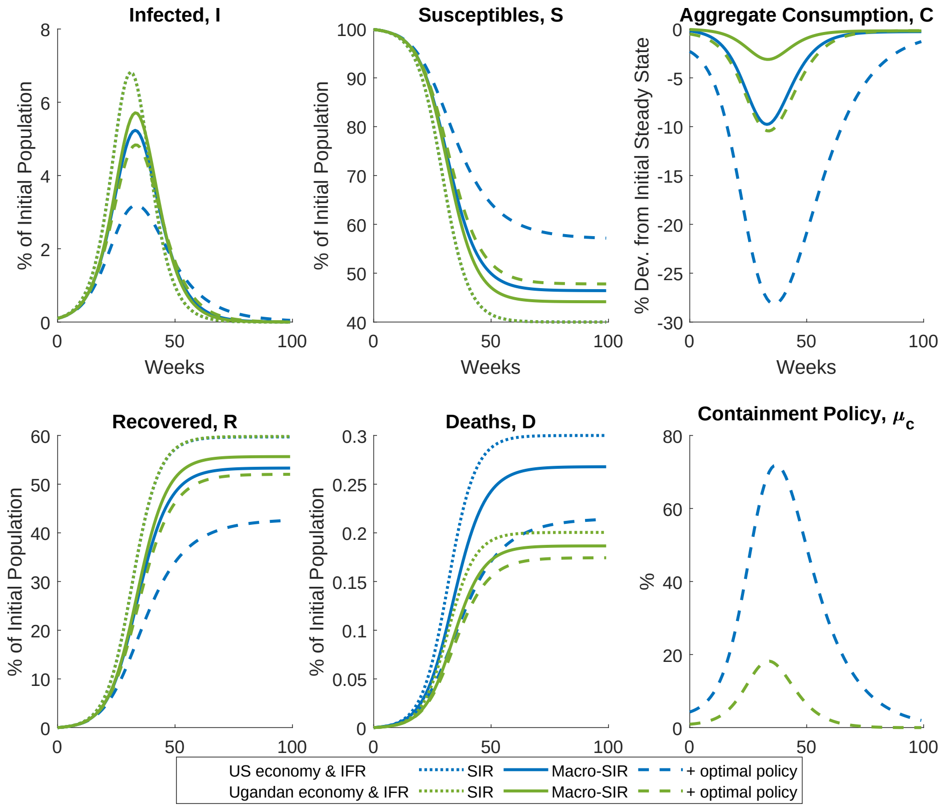

The figure shows the time path of epidemiological (left and middle panels) and economic components (right panels) of the ERT model calibrated to the United States (blue) and Uganda (green). Dotted lines report results from the basic SIR model. Solid lines include agents' voluntary adjustments. Dashed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

reductions lead to relatively large losses in utility, and it becomes less optimal to avert deaths.

This exercise highlights that despite equal death risk, the social planner would choose less stringent containment measures in a poorer economy, where reduction in consumption is more costly in utility terms, i.e., an additional consumption reduction takes a relatively larger share of agents' utility.

3.3 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. The results are presented in Figure 5. Faced with a relatively low death rate, agents themselves find it rational to reduce consumption only marginally, reducing the death rate by 6.9% only. Optimal policy now peaks at 18% and reduces deaths by only an additional 6.1%. This difference to the US setting would be even starker if we used our estimates of the IFR for the US, which is 0.2 percentage points higher.

Figure 6: Low income calibration with low income IFR and subsistence constraints

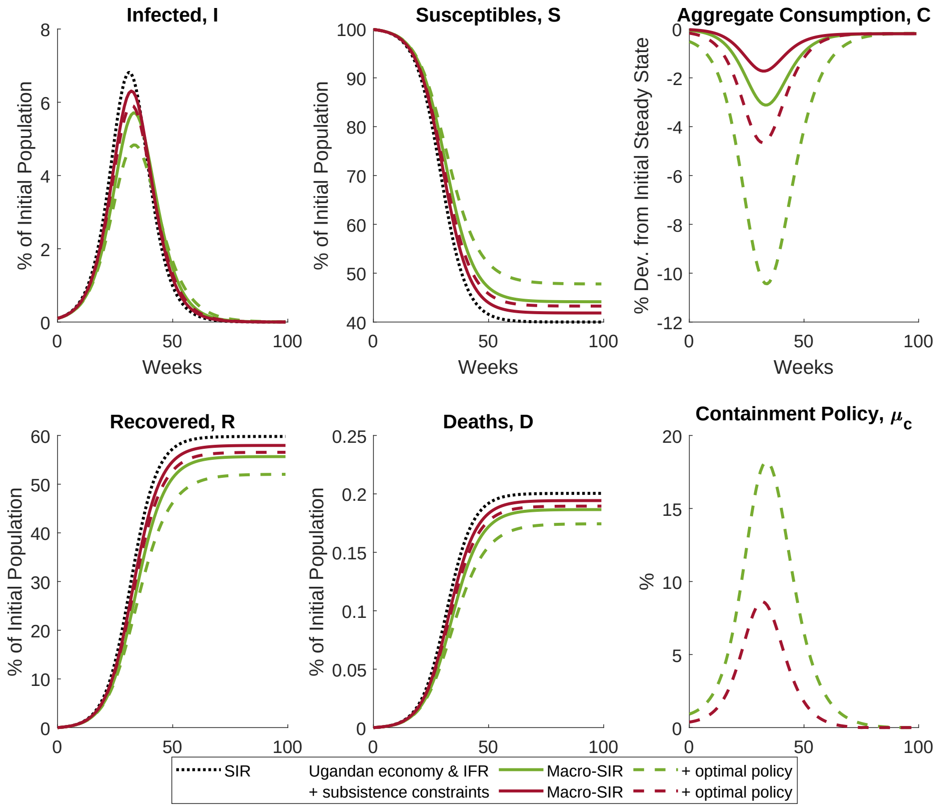

The figure shows the time path of epidemiological (left and middle panels) and economic components (right panels) of the ERT model calibrated to Uganda with (green) and without (red) subsistence constraints. Dotted black line reports results from the basic SIR model. Solid lines include agents' voluntary adjustments. Dashed lines show results from social planner problem maximizing overall utility.

3.4 The role of subsistence constraints

The third step of the calibration to Uganda is to introduce a subsistence constraint into the utility function of the agent, which then becomes

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

We set the subsistence level at \$200, in line with the median monthly nominal wage for female workers without any formal education.

Introducing this constraint makes it even more costly for people with a low income to reduce their labor hours. In case of the United States, the introduction of such a subsistence level would not have any substantial effect as consumption is a lot larger than the threshold in the typical household. In Uganda however, the agents are relatively closer to the constraint and thus the introduction of a subsistence level of consumption in the model makes hours and consumption adjust even less in

Uganda as a response to an increased tax. In total, deaths are only reduced by 5.5% relative to the pure SIR model.

While our calibration implies that agents are still well above the subsistence constraint, and we hence observe a small consumption adjustment, this extension would in principle allow for modeling of agents being pushed against their subsistence constraint by a containment policy.

3.5 The role of beliefs about transmission risk in the acceptance of lockdown policies

The previous subsections have explored optimal policy for overall utility from the perspective of a recent model. We found that for a calibration to Uganda, voluntary and even optimal adjustments are quite small. This stands in contrast to the strict lockdown policies imposed in many developing countries, among them Uganda. Given the large difference between optimal and observed policies, it may be surprising that recent surveys in Senegal and Pakistan have found broad agreement with the measures mandated by the government, despite households reporting substantial reductions in income. A clear majority of respondents (70%) of a nationally representative survey conducted in Senegal in the second week of April supported a then two-week lockdown to curb the spread of the epidemic (Moscoviz and Le Nestour, 2020). Interestingly, this number was similar among the people who had already seen income losses, and higher among those stating to be more worried about the epidemic in general. At the time, the Senegalese government had already imposed a nightly curfew and closed public spaces, and more households than usually had cut down on the size of meals during the survey period. In a non-representative but geographically broad sample from Bangladesh (Brac, 2020), respondents similarly reported large decreases in income of up to 75% due to the drying up of casual labor markets, while at the same time generally supporting further restrictive measures.

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 US and Sweden presented in Farboodi et al. (2020). In particular, this may be the case if agents overestimate the risk of getting infected and/or the risk 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. With a small extension, the ERT model lends itself to an analysis of the necessary overestimation.

In particular, we note that the 'true' infection risk from economic activity $(\pi_{s1}, \pi_{s2}, \pi_{s3})$ are, in reality, unlikely to be known by the agents. It appears more probable that beliefs on both the risk of

Figure 7: Optimal policy with inaccurate beliefs about transmission probabilities

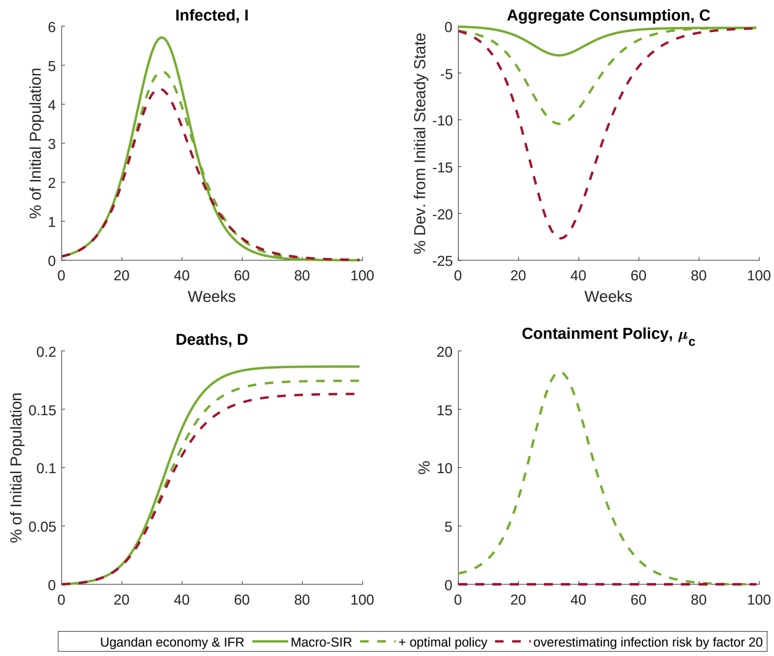

The figure shows the time path of epidemiological (left panels) and economic components (right panels) of the ERT model calibrated to Uganda (green). Solid line includes agents' voluntary adjustments. Dashed green line shows results from social planner problem maximizing overall utility. Dashed red line shows results when agents overestimate infection risk by factor 20.

contracting an infection and of dying from it, at least during early stages of the epidemic, often largely overestimate the true (unknown) parameters in both developing and developed countries. Within the model, we can thus introduce a factor ρ by which agents overestimate the true infection parameters from consumption or work. A similar exercise could be performed for the probability of death. This transforms Equation 2 into

$$\begin{aligned} U_t^s &= u\left(c_t^s, n_t^s\right) + \beta\left[\left(1-\tau_t\right)U_{t+1}^s + \tau_t U_{t+1}^i\right] \\ \text{with } \tau_t &= \rho\pi_{s1}c_t^s\left(I_t C_t^I\right) + \rho\pi_{s2}n_t^s\left(I_t N_t^I\right) + \pi_{s3}I_t \end{aligned}$$

In our calibration of the ERT model to Uganda using the Ugandan IFR, agents reduce consumption voluntarily by 1%, and the optimal policy increases this reduction to 4%. If, however, agents overestimate the infection risk from consuming and working by factor 20, they voluntarily reduce consumption by 10%, 10 times their original voluntary adjustment. This simple exercise illustrates that if agents are overly afraid of contracting the virus (or similarly, overestimate the IFR), their consumption reduction may exceed what the social planner mandates. In turn, this suggests 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. The exercise also suggests that governments may find it harder to restrict economic activity once people get more accurate or even too optimistic beliefs about the medical risks. In this light, early and strict lockdowns may be effective in making containment easier later, as individuals perceive the danger as large and imminent.

4 Conclusion

We apply a standard SIR epidemiological model, and the recent extension with behavioral responses introduced by Eichenbaum et al. (2020), to compare optimal policy responses to the current pandemic for two countries, the United States and Uganda. We calibrate the model with country-specific distributions of age, comorbidities, and income, and extend the standard framework to allow for a subsistence constraint on consumption and a 'fear parameter' that may induce agents to overshoot in their adjustment given the true risks associated with the disease. We show that the optimal policy response to the pandemic, maximizing overall utility in our framework, is much more modest for Uganda than for the United States. The main reason is that reductions of consumption are more costly for poorer agents in general, and in particular for those close to the subsistence constraint. Furthermore, given that the demography of Uganda compared to the United States results in a substantially lower likelihood of a fatal outcome given infection, the health damage in a poorer country is predicted to be lower.

These very different optimal policy responses stand in contrast to the relatively similar *actual* policy responses that we see across the world. Our results suggest that actual strict 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 the 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 explanation, we explored the role of possible overestimation of the fatality risk among the population, leading to demand for stronger measures.

Finally, we reiterate that a central assumption in this exercise is that the disutility of death is equal to the foregone utility of living. While this assumption is 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

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

Alvarez, F., Argente, D., Lippi, F., 2020. A Simple Planning Problem for COVID-19 lockdown. Working Paper .

Atkeson, A., 2020. What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios. NBER Working Paper Series , 25URL: `http://www.nber.org/papers/w26867`.

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.

Barnett-Howell, Z., Mobarak, A.M., 2020. The Benefits and Costs of Social Distancing in Rich and Poor Countries. Working Paper URL: `http://arxiv.org/abs/2004.04867`, `arXiv:2004.04867`.

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 URL: `http://www.hdr.undp.org/sites/default/files/charmes_hdr_2015_final.pdf`.

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. CEPR Covid Economics 15.

Eichenbaum, M., Rebelo, S., Trabandt, M., 2020. The Macroeconomics of Pandemics. NBER Working Paper 26882. doi:10.1017/CBO9781107415324.004, `arXiv:arXiv:1011.1669v3`.

Farboodi, M., Jarosch, G., Shimer, R., 2020. Internal and External Effects of Social Distancing in a Pandemic. Working Paper .

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

Ghisolfi, S., Almås, I., Sandefur, J., Carnap, T.V., Bold, T., 2020. Calculating regional differences in COVID-19 fatality rates. Working Paper .

Glover, A., Heathcote, J., Krueger, D., Ríos-Rull, J.V., 2020. Health versus Wealth: On the Distributional Effects of Controlling a Pandemic * , 0–37.

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

Kermack, W., McKendrick, A., 1927. A Contribution to the Mathematical Theory of Epidemics. Proceedings of the Royal Society 115. doi:10.2307/2304150.

Krueger, D., Uhlig, H., Xie, T., 2020. Macroeconomic dynamics and reallocation in an epidemic. CEPR Covid Economics .

Mahoney, N., 2020. Don't fall for the false trade-offs of COVID-19 policy. URL: https://review.chicagobooth.edu/economics/2020/article/don-t-fall-false-trade-offs-covid-19-policy.

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.

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. Working Paper .

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

Summers, L., 2020. Trump is missing the big picture on the economy. URL: https://www.washingtonpost.com/opinions/2020/03/24/trump-is-missing-big-picture-economy/.

Toxvaerd, F., 2020. Equilibrium social distancing. Cambridge Working Papers in Economics 2021.