How Should Policy Responses to the COVID-19 Pandemic Differ in the Developing World?[*]

Titan Alon   Minki Kim   David Lagakos   Mitchell VanVuren

May 21, 2020

Abstract

The COVID-19 pandemic has already led to dramatic policy responses in most advanced economies, and in particular sustained lockdowns matched with sizable transfers to much of the workforce. This paper provides a preliminary quantitative analysis of how aggregate policy responses should differ in developing countries. To do so we build an incomplete-markets macroeconomic model with epidemiological dynamics that features several of the main economic and demographic distinctions between advanced and developing economies relevant for the pandemic. We focus in particular on differences in population structure, fiscal capacity, healthcare capacity, the prevalence of "hand-to-mouth" households, and the size of the informal sector. The model predicts that blanket lockdowns are generally less effective in developing countries at reducing the welfare costs of the pandemic, saving fewer lives per unit of lost GDP. Age-specific lockdown policies, on the other hand, may be even more potent in developing countries, saving more lives per unit of lost output than in advanced economies.

1. Introduction

Governments in most advanced economies have responded to the COVID-19 pandemic with dramatic and unprecedented policy responses. Huge swaths of the economy have been ordered shut down and millions of workers required to stay home indefinitely. To cushion against the loss of income, governments have sent direct transfers to workers affected by the lockdowns, in addition to making regular social insurance payments. Unemployment benefits, in particular, have reached levels many times larger than at any prior point in history. The total costs of these and other transfers may even reach 10 percent of GDP in some countries.

As COVID-19 made its way to less-developed countries, policy makers there largely followed suit with similarly sweeping lockdowns. Yet it quickly became clear that policy responses in the developing world could not just mimic those of the west. Policy would instead have to be tailored to the dramatically different economic and demographic landscape that characterizes most low-income countries. Governments in these countries lack the fiscal capacity to make substantial transfers to major segments of the population for long periods. Moreover, high levels of poverty mean that many households are effectively living hand-to-mouth, rendering prolonged lockdowns economically infeasible. Large informal sectors make lockdown enforcement much harder, and expanding the tax base nearly impossible. The potential health consequences of the pandemic are also quite distinct from those of the west, which has an older and more susceptible population but substantially more developed healthcare systems.

This paper provides a preliminary quantitative analysis of how policy responses to the COVID-19 pandemic should differ in developing countries. We focus on broad aggregate economic policy, meaning over the extent and duration of lockdowns and transfers, including the extent to which either should be age-specific. We also illustrate how the fiscal and healthcare capacity constraints interact with informality and the younger demographics to inform optimal policy responses relative to advanced economies of the west. The analysis is preliminary largely because the requisite data are not available yet to draw firmer conclusions, though data is being collected in developing nations at an unprecedented rate, which will soon better inform analyses such as the current one.

To study how aggregate policy should differ in developing and advanced economies, this paper follows the newly emerged literature on the macroeconomics of pandemics by combining a workhorse macro model with a variant of the SIR model standard in epidemiology. The model in this paper builds most closely on the framework of Glover, Heathcote, Krueger, and Ríos-Rull (2020), and in particular their tractable model of heterogeneous agents that face income shocks and health risk that varies by age. In addition, our model features fiscal capacity constraints –

motivated broadly by the work of Besley and Persson (2009), Jensen (2019) and many others – which reduce the ability of governments to tax and transfer resources across households. It adds an "informal sector" as in Ulyssea (2018) and a large development literature, and worker sorting on skill as in Roy (1951) and following the specification of Feng, Lagakos, and Rauch (2018). We allow for hand-to-mouth consumers, which is the emphasis of Kaplan, Moll, and Violante (2020), and healthcare capacity constraints as in the recent macroeconomic literature on the pandemic. We model the epidemiological dynamics in the model using a variant of the SICR model that is standard in the epidemiology literature, combined with economic choices that govern the disease path endogenously, as in the macroeconomic literature on disease transmission (e.g. Greenwood, Kircher, Santos, and Tertilt, 2019; Eichenbaum, Rebelo, and Trabandt, 2020; Guerrieri, Lorenzoni, Straub, and Werning, 2020; Alvarez, Argente, and Lippi, 2020; Chang and Velasco, 2020).

We parameterize the model to match the pre-pandemic stationary distribution of a representative advanced economy, calibrated to match characteristics of countries in the top quartile of the world income distribution. We then compute the model's equilibrium response to the COVID-19 pandemic as a surprise "MIT shock" where a small exogenous fraction of the population becomes infected with the coronavirus, and the disease then makes its way through the populous. We then do the same for an alternative calibration of the model taken to match a representative developing economy, representing averages of the lowest-income quartile of countries. We match in particular the substantially lower population share of the old, the lower fiscal and healthcare capacity, the larger informal sector, and greater proportion of hand-to-mouth households. We then simulate the effects of various types of lockdowns in the advanced and developing economies and compare their impacts on GDP, fatalities, and consumption-equivalent welfare.

The model predicts that blanket lockdowns (affecting the entire population) are not as effective in developing countries as in advanced countries. In particular, blanket lockdowns do less to reduce the welfare costs of the COVID-19 pandemic in developing countries, and save fewer lives per unit of lost GDP. For example, in the developing economy, a medium length blanket lockdown lasting 28 weeks saves around 70 lives per hundred thousand at the cost of a 7 percent decline in GDP. Thus, for every unit of GDP lost, the policy saves 10 lives per hundred thousand people. The same length blanket lockdown in the advanced economy reduces GDP by 16 percent and saves around 320 lives per hundred thousand people, amounting to 20 lives saved per hundred thousand for every lost unit of GDP. By this metric, developing countries save about half as many lives per unit of economic output lost as advanced economies. Using a consumption equivalent welfare metric, the lockdown in the advanced economy raises wel-

fare by around 3 percent, compared to just 0.6 percent in the developing economy. We find that lockdowns that are shorter or longer are less effective, though still better than having no lockdown at all.

Blanket lockdowns in our model have sharply differing impacts across young and old households, just as in the work of Glover et al. (2020), Bairoliya and Imrohoroglu (2020), Acemoglu, Chernozhukov, Werning, and Whinston (2020) and others. Older households gain a lot more from lockdowns, since they have the greatest reduction in fatality risk. We find, in fact, that these heterogeneous impacts are even starker in developing economies, pointing to a potential role for age-specific lockdowns there as well. We therefore simulate the role of age-specific lockdowns that require that just older individuals remain in lockdown, and with larger transfers to each old individual such that the total amount spent on transfers is the same as under the blanket lockdown policies studied above.

Our model predicts that age-specific policies are even more potent in developing countries than in advanced economies. A medium-length age-specific lockdown saves around 95 lives per hundred thousand for every lost unit of GDP, which is twice as much as in advanced economies, and ten times as much as under blanket lockdowns in developing economies. The reason is that age-specific policies allow governments to isolate only those with the highest fatality risk, and to provide them with larger transfers than under blanket lockdowns. This is particularly attractive in developing countries, since older individuals reflect such a small share of the total population there. Overall, the quantitative analysis in this paper points to age-specific lockdowns as the most promising form of lockdown for developing countries, though of course many logistical issues are still open.

There are many aspects of developing economies that we have not modeled in this paper, and many of these are surely relevant for the study of optimal responses to the COVID-19 pandemic. Differential testing and tracing policies (Berger, Herkenhoff, and Mongey, 2020) are absent here but surely worth studying. Our model also abstracts from differential impacts by gender (Alon, Doepke, Olmstead-Rumsey, and Tertilt, 2020), policy uncertainty (Baker, Bloom, Davis, and Terry, 2020), and stock-market impacts (Toda, 2020), as well as differences in sanitation, living conditions and co-morbidities that may affect fatality rates. Perhaps the most conspicuous absences are open-economy considerations, given that developing countries are already witnessing negative impacts from capital outflows, and the effects of the decline in natural resource prices, which further depress fiscal space. We plan to add both of these to our analysis in the future, and other papers are already doing the same, in particular Çakmaklı et al. (2020), who study the effects of COVID-19 in Turkey.

## 2. Motivating Facts for Differing Policy Responses in Developing Countries

We begin by summarizing four important demographic and economic characteristics that differ across countries in ways that we view as essential for understanding how policy responses to COVID-19 should differ in developing economies. These are: (i) the much younger populations in developing countries, (ii) their lower fiscal capacity, (iii) their more widespread informality, and (iv) their lower healthcare capacity. To be sure, all of these patterns are already known in some form or another. The goal of this section is to present these patterns in a systematic way and briefly highlight their relevance for the effects of the pandemic, which will help motivate the model and quantitative analysis that follows.

2.1. Younger Populations

All available evidence so far suggests that COVID-19 poses dramatically different health risks to older individuals, in particular those over the age of 65. Early centers of infection in the west, such as Italy, experienced health impacts concentrated on those in this older age range, with particularly severe fatality rates for those in their 80s and 90s. At the same time, the number of deaths linked to COVID-19 for those under 20 has been negligible, though certainly not zero.

A basic demographic difference between advanced and developing economies is that populations are far younger in the developing world. Since fatality rates from COVID-19 are very low for young individuals but rise sharply with age, these demographic differences suggest much smaller populations of vulnerable individuals in the developing world. One can see these demographic differences starkly when looking at cross-country data on the median age. Figure 1 plots the median age against GDP per capita in a set of 158 countries using data from UN Population Division and Penn World Tables. Data from the UN Population Division show that countries in the bottom quartile of the world income distribution have a median age of 19.1 years. Nigeria, Africa's most populous country, has a median age of 17.9, while countries like Angola and the Democratic Republic of the Congo have median ages of just 16.4 and 16.8 years old. By contrast richer countries like Italy, the United Kingdom and France have median ages of 45.9, 40.2 and 41.2, respectively.

Figure 1: Median Age of the Population

Note: This figure plots the median age in 2015 across 158 countries. Median age is defined as the age that divides the population in two parts of equal size, that is, there are as many persons with ages above the median as there are with ages below the median. GDP per capita is expressed at PPP and taken from Penn World Table 9.1 (Feenstra et al., 2015). Median age data is from the UN Population Division.

Another statistic indicative of the much smaller vulnerable population in the developing world is the cross-country data on the population above 65. Figure A.1 plots the fraction of the population that is above 65 against GDP per capita in a set of 162 countries using data from the World Bank and the Penn World Tables. In the world's poorest countries the fraction of the population that is above age 65 is negligible, with an average of around 3 percent for countries in the bottom quartile of the world income distribution. The older population is much larger as a fraction of the total in richer economies, and reaches around one quarter of the population in Japan. Among countries in the topic quartile of the world, the average is about 15 percent of the population being above age 65.

It is hard to look at statistics like these and not see how sharply different the impacts of COVID-19 will be in less developed countries. Concretely, while almost everything about COVID-19

suggests a more severe impact in less-developed countries, the far younger demographic is clearly in their favor. In the model that we present in the following section, we reserve a central role these sharp demographic differences, and we explore how important age differences are for optimal policy responses to the pandemic.

2.2. Lower Fiscal Capacity

Developed nations take for granted the ability for their governments to raise tax revenues and use the proceeds to provide public goods and make transfers. This fiscal capacity is not shared by the public sectors in developing countries, as a long literature has emphasized (see Besley and Persson, 2013, for an overview). This literature has emphasized how developing nations generally lack the ability to monitor and enforce tax payments from its citizens, and have less efficient revenue authorities than do richer countries.

As a crude, but widely used, measure of fiscal capacity across countries, Figure 2 plots total tax revenues as a fraction of GDP taken from the ICTD Government Revenue Dataset against GDP per capita. The nearly linear positive relationship between taxation relative to GDP and income per capita highlights the much lower role that taxation plays in less developed economies. Similar patterns have been observed in the time series for countries as they develop and increase rates of taxation, particularly on labor income (Besley and Persson, 2013; Jensen, 2019). Although these patterns don't prove that developing economies have larger hurdles in raising tax revenue and spending public funds effectively – as opposed to just choosing to tax less – they are certainly consistent with the interpretation of lower fiscal capacity given by the literature, and taken in this paper.

The lower fiscal capacity in poorer countries is relevant for studies of the pandemic for several reasons. Most importantly, it limits the ability of governments to institute large-scale income replacement programs for furloughed workers during lengthy lockdowns or in response to widespread business closures. Not receiving any payments is a clear disincentive for citizens to comply with a lockdown, especially for those that have little savings to fall back on. In addition, the inability to raise taxes effectively limits governments' ability to borrow, which further reduces their ability to make payments to furloughed workers. Other types of fiscal policy, such as Keynesian stimulus spending, are also limited by low fiscal capacity.

2.3. Larger Informal Sectors

A large share of employment in developing countries is concentrated in informal production activities. By definition, such markets are beyond the purview of the state to tax or regulate, and make law enforcement difficult. Lockdown policies, which call for citizens not to leave

Figure 2: Fiscal Capacity, Proxied by the Ratio of Taxes to GDP

Note: This figure plots the ratio of taxation to GDP ratio across 116 countries. Tax revenues including social contributions. GDP per capita is expressed at PPP and taken from Penn World Table 9.1 (Feenstra et al., 2015). Tax revenue data is from the ICTD / UNU-WIDER Government Revenue Dataset 2019 (ICTD/UNU-WIDER, 2019).

home for work or for other public interactions, are difficult to enforce anywhere. Clearly, enforcing lockdowns is even harder in an environment with only a minority of the workforce employed at formal businesses.

While few would disagree that informality is more prevalent in less developed economies, measuring informality is not a straightforward enterprise. Figure 3 plots the size of the informal sector as a fraction of non-agricultural employment, as measured by the International Labor Organization (ILO), against income per capita. The ILO defines informal workers as those that produce goods or services meant for sale or barter. Self-employed street vendors, taxi drivers and home-based workers, regardless of size, are all considered informal workers. Excluded are agricultural and related activities, as are any households producing goods exclusively for their own use (e.g. subsistence farming, domestic housework, care work, and employment of paid

Figure 3: Size of the Informal Sector
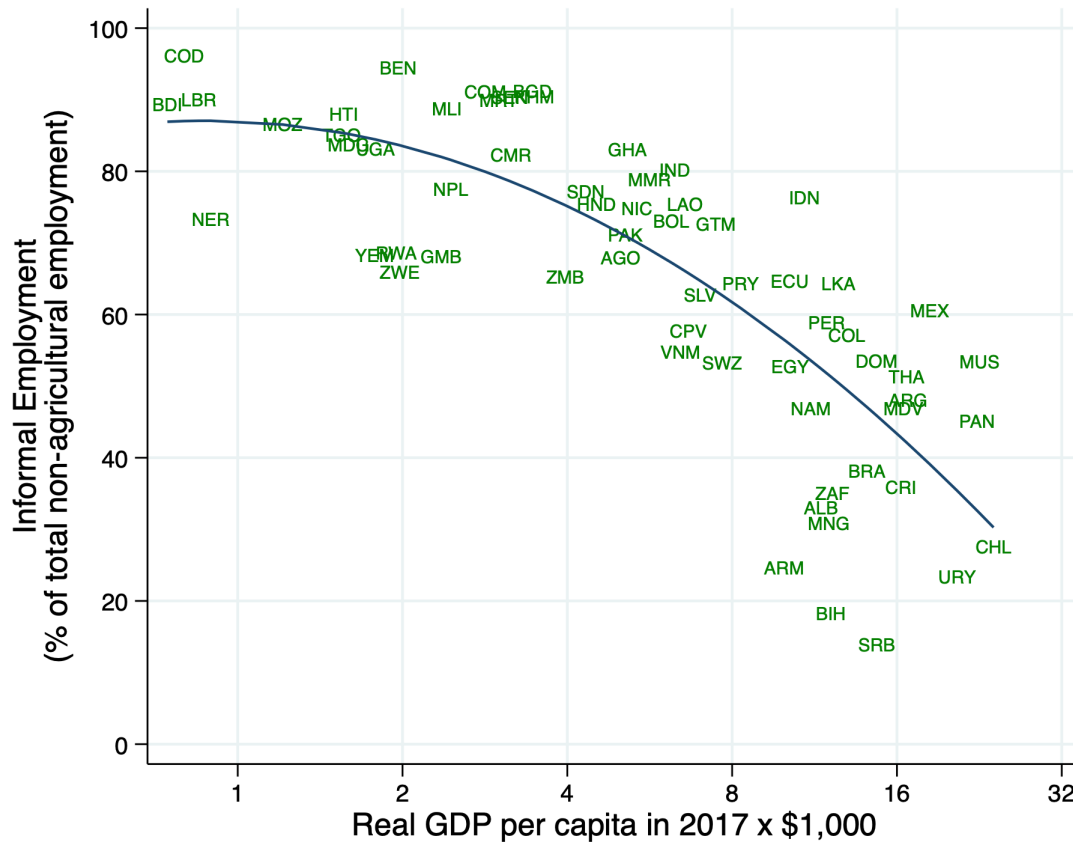
Note: This figure plots the employment in the informal economy, measured by the ILO, as a percentage of total non-agricultural employment, in 63 countries. The informal economy is defined as in the text. GDP per capita is expressed at PPP and is taken from the Penn World Table 9.1 (Feenstra et al., 2015). Informality data is from the ILOSTAT database.

domestic workers), and volunteer services.

Figure 3 shows a sharp decline in informality rates in non-agricultural activities with GDP per capita. The countries with the lowest income have informality rates above 80 percent in most cases. The richest countries in this ILO database have informality rates below half in most cases, though this figure excludes most of the richest countries in the world, which undoubtedly have even lower rates of informality. A related statistic, for which data is readily available for almost every country in the world, is the fraction of the workforce that is self-employed. This fraction, which we plot in Appendix Figure A.2, runs from close to the entire workforce in the poorest countries to virtually none of it in the richest countries. This well-known pattern reinforces the fact that employment takes on a very different form in poorer countries, with own-account workers and family businesses being the dominant source of labor inputs.

An obvious way in which informality ties the hands of governments in low-income economies is that it reduces their ability to collect additional taxes and make transfers to those in lockdown. The widespread informality and limited fiscal capacity are of course very closely linked, with each reinforcing the other. During the pandemic, any attempts to keep households in lockdown may result in increases in the size of the informal sector, which may hurt attempts to control disease or raise new revenues.

A related feature of the informal sector relevant for policy responses to the pandemic is the concentration of low-skilled jobs there. To the extent that lockdown policy forces desperate workers that have run down their savings to enter the informal sector, these workers can be expected to perform marginal tasks that do not generate much income. For workers already in marginal tasks in the formal sector, this may not represent much of a change. But for those in more skilled jobs to begin with, having to work in low-skilled informal activities would represent a substantial decline in household income and therefore consumption. If enough workers become desperate and enter the informal sector, this could reduce aggregate productivity and further shrink the government's tax base.

2.4. Lower Healthcare Capacity

Developing countries typically have substantially less ability to control disease than do richer countries. Sanitation and hygiene are more of an issue given the lack of widespread piped water and functioning sewage systems. Health infrastructure, especially hospital and health clinic capacity, is also less developed. For mild cases of COVID-19 infections, this may make little differences, as bed rest is likely to suffice in these mild cases. However, for critical cases, the lack of intensive-care capacity is a clear disadvantage for developing countries in their attempts to save lives during the pandemic.

Figure 4 plots the number of hospital beds per 10,000 people, as reported by the World Health Organization (WHO), against GDP per capita. The number of hospital beds is an imperfect measure of hospital capacity for many reasons, most importantly because it is not a bed per se that helps critical patients recover from COVID-19 but trained doctors, equipment like ventilators, and appropriate pharmaceuticals. Still, for lack of more comprehensive cross-country data, we take hospital beds as a proxy for medical care capacity.

By this metric there are stark differences in healthcare capacity across countries. Richer countries, which have quite some range amongst themselves, average around 49 hospital beds per 10,000 people. Countries like Japan and Korea have even more beds per capita, having 134 and 115 beds per 10,000 people, respectively. This is still far higher than the capacity in developing countries, which is a paltry 12 beds per 10,000 people on average in the bottom quartile of the

Figure 4: Healthcare Capacity, Proxied by Hospital Beds per 10,000 People

Note: This figure plots the number of hospital beds available per 10,000 inhabitants in 153 countries. GDP per capita is at PPP and taken from the Penn World Table 9.1 (Feenstra et al., 2015). The hospital bed data are from the World Health Organization's Global Health Observatory.

income distribution. In Appendix Table B.2, we report the availability of intensive care unit (ICU) beds and per capita healthcare costs across a limited set of countries. Consistent with the patterns observed from the number of hospital beds, it appears that low income countries possess significantly fewer ICU beds than high income countries. Systematic data on ventilators are harder to come by, but the available evidence so far points to even starker differences in ventilator supplies across countries. According to the New York Times, there are fewer than 2,000 working ventilators across 41 African countries, as of April 18, 2020. South Sudan has four ventilators for a population of 11 million, the Central African Republic has three for a population of five million, to name a few. Ten countries in Africa have none at all.[1]

---

[1]See Maclean, Ruth and Marks, Simon, "10 African Countries Have No Ventilators. That's Only Part of the Problem", April 18, 2020, The New York Times

3. The Model

Our analysis draws on a quantitative heterogenous-agent macroeconomic model with epidemiology as in the SICR model to analyze how policy responses to the COVID-19 pandemic should differ in developing countries. The model is equipped with several features that vary between advanced and developing economies that are relevant for the pandemic response, as motivated by the data presented in the previous section. These include uninsurable idiosyncratic health and income risks, age heterogeneity, fiscal capacity constraints, an informal sector, and healthcare capacity constraints. This section now presents these features in detail.

3.1. Households and Preferences

The economy is populated by a unit mass of heterogenous households who make consumption, savings, and sectoral employment decisions subject to idiosyncratic income and health risks. Individuals differ in their age $j \in \{\text{young}, \text{old}\}$, initial assets a, and permanent labor productivity $z \sim G$. Time is discrete and each period represents two weeks. Household preferences are given by:

$$U = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta_j^t \left\{\log(c_t) + \bar{u}\right\}\right], \tag{1}$$

where $\beta_{\text{young}} < \beta_{\text{old}}$ captures age heterogeneity in the population and $\bar{u}$ represents the flow utility value of being alive. We follow the tractable formulation of Glover et al. (2020) which abstracts from explicitly modeling age, appealing to the logic that pandemics are sufficiently short-lived relative to entire lifetimes. It thus suffices to model only the expected number of years left to live, which is captured by the heterogeneity in discount factors. The term $\bar{u}$ follows the specification of Jones and Klenow (2016), and represents the reason that model households try to avoid fatality risk.

Households can choose to work in either the formal ($s = f$) or informal ($s = i$) labor markets where they can earn wage w_s per effective hour worked. At the beginning of life, workers draw their permanent sectoral productivity, $z \sim G$, and choose occupations as in a Roy (1951) model with one-sided selection. Since work in the informal sector is largely unskilled, we normalize z to unity for this sector so that there is no within sector variation in permanent productivity, as in the specification of Lagakos, Mobarak, and Waugh (2019).

Incomes in both sectors are subject to idiosyncratic productivity shocks as in the Aiyagari-Bewley-Huggett framework (Bewley, 1977; Huggett, 1993; Aiyagari, 1994). Specifically, we assume that individual labor productivity in each sector is composed of the sector-specific per-

manent component z and an idiosyncratic component ν following the stochastic process,

$$\log \nu_{t+1} = \rho_\nu \log \nu_t + \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim F(0, \sigma_\nu). \tag{2}$$

We include idiosyncratic income risk because developing countries are far from having full insurance, and so accounting for how people insure themselves in response to policies which may keep them away from work for prolonged periods of time is a first order consideration.

After choosing their occupation and observing their income realization, households make consumption and savings decisions given the interest rate, r, and subject to a no-borrowing condition, $a \geq 0$. Formally, the household budget constraint is given by,

$$c + a' \leq \mathbb{1}_{\{s=i\}} w_i \nu + (1-\tau) \times \lambda_{LD}^w \times \mathbb{1}_{\{s=f\}} w_f z \nu + (1+r)a + T \tag{3}$$

where τ is the income tax rate and T is government transfers. The term λ_{LD}^w parameterizes productivity lost during the imposition of a government lockdown and is equal to one in normal times and equal to $\lambda_{LD}^w = \lambda_w < 1$ during lockdowns. Importantly, limited state capacity implies that government taxes and commercial restrictions, such as the lockdown, can only be enforced in the formal employment sector. In reality, enforcement capabilities are probably more nuanced, but it is almost certainly much easier in formal places of business than in informal activities. The possibility of moving into the informal sector in response to a lockdown is similar to the movements out of market activities at the start of the pandemic emphasized in Krueger, Uhlig, and Xie (2020), and broadly consistent with the evidence of Zhao, Storesletten, and Zilibotti (2019) that workers respond to economic downturns by moving back into agriculture. There is substantial evidence since the onset of the pandemic in the developing world that many workers do indeed respond by moving into rural agriculture.

3.2. Aggregate Production Technology

The economy produces a single final good by combining domestic and foreign capital with labor services supplied by the formal and informal employment sectors. The aggregate production technology is given by,

$$Y = L^\alpha K^{1-\alpha},$$

where $0 < \alpha \leq 1$ is labor's share of value-added. The aggregate capital stock is composed of both domestic and foreign sources, $K = K^D + K^F$, which can be rented at an exogenously given international rental rate r^F (different from r, as we explain below) and which depreciates at rate δ.

Aggregate labor depends on the supply of both formal and informal labor services. Since skilled work is largely concentrated in the formal sector, and unskilled work in the informal sector, it is natural to model the two labor inputs as gross-substitutes in the aggregate (as in Ulyssea, 2018). Formally, aggregate labor supply is given by,

$$L = \left[A L_f^{\frac{\sigma-1}{\sigma}} + L_i^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

where $0 < \sigma < \infty$ is the elasticity of substitution between formal and informal labor services and A indexes the relative productivity of formal sector employment. We allow technology A to augment skilled labor, and not unskilled labor, since a large literature finds that cross-country productivity differences are skill-biased, rather than skill neutral (Caselli and Coleman, 2006; Malmberg, 2018). This assumption is important for the sorting of workers by skill in the model, and the prediction that workers with higher skill (permanent productivity) levels sort into the formal sector, with less productive workers selecting into the informal sector.

3.3. Credit and Capital Markets

Credit market incompleteness prevents households from borrowing against future earnings. As a result, individuals must maintain assets $a \geq 0$ in formulating their consumption plans subject to (3), giving rise to hand-to-mouth consumers as well as a precautionary savings motive in response to idiosyncratic health and income risks. The precautionary motive is important for getting aggregate welfare measurements correct since it creates another feedback between the epidemiological and economic dynamics, as individuals withhold some consumption to increase precautionary savings in response to the pandemic's onset.

Furthermore, financial frictions in capital markets create a spread between the economy's borrowing and savings rates. Specifically, the interest rate paid on domestic savings is such that $r = r^F - \chi$ where $\chi > 0$ represents a financial wedge leading the returns on savings to be less than the rental rate of capital faced by governments and firms borrowing in international capital markets. These frictions increase the number of economically vulnerable hand-to-mouth consumers in developing countries relative to advanced ones, raise the cost of government borrowing to support welfare programs, and distort capital accumulation. This gives us a tractable way of controlling the level of hand-to-mouth consumers in the model, without having to model illiquid assets explicitly as in Kaplan et al. (2020).

3.4. Public Health and Hospital Capacity

Households face idiosyncratic health risk which can reduce their labor productivity and increase the probability of dying. Susceptibility to infection is determined in part by economic decisions taken by households. Once infected, progression of the disease depends on an individual's age and the availability of public health infrastructure offering treatments.

Health risks are modeled using an SICR epidemiological model with five health states: susceptible ($\mathbb{S}$), infected ($\mathbb{I}$), crital ($\mathbb{C}$), recovered ($\mathbb{R}$), and deceased ($\mathbb{D}$). We denote by N_t^x the mass of individuals in each health state $x \in \{\mathbb{S}, \mathbb{I}, \mathbb{C}, \mathbb{R}, \mathbb{D}\}$ at time t and use $N_t = N_t^{\mathbb{S}} + N_t^{\mathbb{I}} + N_t^{\mathbb{C}} + N_t^{\mathbb{R}}$ to measure the non-deceased population.

$$
\mathbb{S} \xrightarrow[\text{informal: } \pi_i^{\mathbb{I}}]{\text{formal: } \pi_f^{\mathbb{I}}} \mathbb{I}
$$

young: $1-\pi_y^{\mathbb{C}}$
old: $1-\pi_o^{\mathbb{C}}$
$\longrightarrow \mathbb{R}$

young: $\pi_y^{\mathbb{C}}$
old: $\pi_o^{\mathbb{C}}$
$\longrightarrow \mathbb{C}$

young: $1-\pi_{yt}^{\mathbb{D}}(N_t^{\mathbb{C}}, \Theta)$
old: $1-\pi_{ot}^{\mathbb{D}}(N_t^{\mathbb{C}}, \Theta)$
$\longrightarrow \mathbb{R}$

young: $\pi_{yt}^{\mathbb{D}}(N_t^{\mathbb{C}}, \Theta)$
old: $\pi_{ot}^{\mathbb{D}}(N_t^{\mathbb{C}}, \Theta)$
$\longrightarrow \mathbb{D}$

Figure 5: Dynamics of Health States and Transition Probabilities

Individuals who contract the virus experience a proportional drop in productivity of $1-\eta$ for one model period (two weeks), at which point they either recover or enter a critical health state. The probability of becoming critically ill depends on an individual's age and is given by π_j^C. Those in critical health are unable to work and require hospitalization. The likelihood of recovery in the hospital depends again on their age in addition to the availability of public health infrastructure, such as ICU beds and ventilators. In particular, the fatality rate of a critically ill patient of age j is given by:

$$
\pi_{jt}^{\mathbb{D}}(N_t^{\mathbb{C}}, \Theta) = \begin{cases} \pi_j^{\mathbb{D}} & \text{if assigned ICU bed} \\ \kappa \times \pi_j^{\mathbb{D}} & \text{if not assigned} \end{cases}
$$

where π_j^D is a baseline fatality rate for age j individuals in critical health and κ governs the impact on fatality rates of strained hospital resources. Whether or not a critically ill patient receives an ICU bed depends on overall hospital capacity and the number of other patients.

Specifically, letting Θ denote hospital ICU capacity, the probability a new patient receives an ICU bed is given by $\min\{\Theta/N_t^C, 1\}$. In other words, all critically-ill patients receive an ICU bed if hospital capacity constraints are not binding, and beds are rationed amongst the critically-ill with probability Θ/N_t^C when constraints bind (Kaplan, Moll, and Violante, 2020).

While the disease's progression is exogenous, the probability a susceptible person becomes infected depends on endogenous economic decisions and the prevalence of infections in the population. Specifically, the baseline probability a susceptible person becomes infected is:

$$\pi^{\mathbb{I}} = \beta^{\mathbb{I}} \times \frac{N^{\mathbb{I}}}{N}$$

where $N^{\mathbb{I}}/N$ is the share of infected people in the population and $\beta^{\mathbb{I}}$ is the "behaviorally adjusted infection rate," which accounts for both the disease's biological transmission rate as well as population wide behavioral responses to avoid being infected (i.e. improved hygiene, social distancing). Using behaviorally adjusted rates is quantitatively important since existing research has shown that these public behavioral responses can substantially reduce transmission rates in practice (Greenwood et al., 2019). Infection rates can be further mitigated by economic lockdowns which shutdown commercial activity. Importantly, such shutdowns only affect the formal sector so that population susceptibility to infection depends in part on individual occupational choice. As a result, individuals working in the informal sector always face the baseline infection rate, while those in the formal sector become infected with probability:

$$\pi_{jt}^{\mathbb{I}}(\text{lockdown}) = \begin{cases} \pi^{\mathbb{I}} & \text{without lockdown} \\ \lambda_h \times \pi^{\mathbb{I}} & \text{with lockdown} \end{cases}$$

where the parameter $\lambda_h \leq 1$ captures the effectiveness of lockdown policies at mitigating the spread of disease.

Our epidemiological model identifies several aggregate health externalities which contribute to the spread of disease. For example, infection probabilities depend in part on the aggregate population of infected individuals. Furthermore, hospitals face congestion inefficiencies which cause the fatality rate to increase as the number of critically ill patients being treated expands. These externalities, and their interaction with economic decisions, creates a margin upon which public health policy, such as lockdowns, could act to improve welfare.

3.5. Government and Taxation

The government has power to tax, transfer, and impose economic lockdowns subject to the constraints imposed by limited fiscal capacity and labor market informality. We further require that the government run a balanced flow budget which satisfies,

$$B_t + \Pi + \tau \int \int \mathbb{1}_{\{s=f\}} y(a,x,v) dQ dG = \Delta \times T$$

where Π represents natural resource revenue or foreign aid, $y(a,x,v)$ is pretax income for individual $(a,x,v) \sim Q$, τ is the prevailing tax rate, and T is aggregate transfers to households. Limited fiscal capacity is captured by the iceberg cost $\Delta > 1$, which require the government raise $1/\Delta$ dollars in revenue for every dollar of transfers to households. Modeling limited capacity through the iceberg costs Δ is a parsimonious way to represent the resources developing countries lose simply trying to collect taxes and the funds that are diverted to self interested parties before being spent on public programs, as emphasized by the literature on fiscal capacity. In the developing world, government resources lost to these inefficiencies is thought to be large, and, in the case of Ghana, have been shown to reach nearly 50 percent of property tax revenue collected (Dzansi et al., 2018).

In addition to tax revenue, we allow developing countries access to emergency bonds, B_t, which can be used to finance additional welfare transfers during government imposed lockdowns. The source of these funds is international donors and multinational institutions such as the IMF, World Bank, and World Health Organization. Funds borrowed for emergency transfers accrue interest at rate $1 + r^F$ until the pandemic ends, at which they are repaid through annual annuities. Formally, emergency transfers are given by:

$$B_t = \begin{cases} \bar{B} & \text{during the lockdown} \\ \frac{r^F}{1+r^F} \times \sum_{t_l - t_s}^{t_l - t_e} \left(1 + r^F\right)^t \bar{B} & \text{after pandemic ends} \\ 0 & \text{otherwise} \end{cases}$$

where $\bar{B}$ is the size of per-period emergency transfers during lockdown, which we take parametrically, and t_s, t_e, and t_l index the lockdown's start, the lockdown's end, and the pandemic's end, respectively.

Alongside its fiscal powers, the government can impose an economic lockdown on the formal sector. While in place, lockdowns reduce disease transmission rates by $1 - \lambda_h$ and reduces pro-

ductivity in the formal sector by $1-\lambda_w$. The pair $0 < \lambda_w, \lambda_h < 1$ can therefore index government lockdown policies, with lower values indicating stricter lockdown measures.

4. Quantitative Analysis

In this section, we discuss the calibration strategy, validate the model's fit, and present the counterfactual results for the benchmark lockdown scenarios and age-dependent policies. After validating the calibration, we simulate the aggregate effect of the COVID-19 pandemic in developing and developed countries with and without aggregate policy responses. Specifically, we study the transition dynamics which emerge when an economy in steady state that is not expecting a pandemic is suddenly hit by the onset of COVID-19 and correctly anticipates its epidemiological dynamics. For each scenario, we report the cumulative changes in welfare, GDP, and aggregate fatalities over the pandemic's duration relative to the pre-pandemic steady state. We also report results on heterogeneity with respect to age and income. The final section reports alternative counter-factual outcomes if existing lockdown measures were made age-dependent, and targeted specifically at helping the elderly population.

4.1. Data Sources and Calibration

For expositional clarity, we divide the calibrated targets into three broad categories corresponding to those governing economic mechanisms, those controlling epidemiological dynamics, and those delineating differences between advanced and developing countries.

Table 1 reports parameters which govern the core economic dynamics of the model. Population demographics are modeled using age dependent discount factors accounting for differences in the remaining years of life for young and old workers. The age specific discount factors are taken from Glover et al. (2020), and the stochastic income processes are taken from Floden and Lindé (2001), who estimate similar income processes in the United States and Sweden. The distribution of permanent productivity $z \sim G$ in the formal sector is modeled by a Fréchet distribution with shape parameter ϕ, taken from Lagakos and Waugh (2013). While our formulation allows for imperfect substitutability between employment sectors, we take these to be perfect substitutes in our initial exercises. Finally, labor's share of income comes from Gollin (2002), and the rental rate of capital is set to the two-week return on pre-COVID Treasury Bills.

Table 2 reports parameters controlling the epidemiological transmission of disease and their interactions with public health infrastructure and lockdown policies. We take parameters governing the age-dependent disease progression probabilities from the epidemiological simulation studies of Ferguson et al. (2020). The effect of hospital congestion on disease fatality

Var	Description	Value	Source / Target
r^F	Exogenous interest rate	0.0006	Pre-COVID T-Bills rate 1.5%
ϕ	Shape-parameter of Frechet distribution G	2.7	Lagakos and Waugh (2013)
ρ_v	Persistence of idiosyncratic income shock	0.91	Floden and Linde (2001)
σ_v	St.Dev of idiosyncratic income shock	0.04	Floden and Linde (2001)
α	Labor share	0.6	Gollin (2002)
β_y	Discount factor for the young	0.9984	Glover et al. (2020)
β_o	Discount factor for the old	0.9960	Glover et al. (2020)

Table 1: Calibration of Economic Parameters

rates, κ, is taken from Glover et al. (2020). The productivity penalty of becoming infected, η, is set to match an 80 percent share of asymptomatic infection cases; such a choice is motivated by the observation that those known to be infected cannot work, and so have productivity of zero, while those who are infected but asymptomatic may continue to work unhindered.

The behavior-adjusted infection generating rate is chosen to match peak rates in an unchecked pandemic. Specifically, in the SIR class of models without policy interventions, there is generally a direct link between the infection generating rate and the peak infected population. We use this relationship to infer transmission rates, taking the expected infection peak from the cross-country estimates of Toda (2020).[2] The advantage of this approach is that while infection generating rates vary with the time-scale of a model, peak infection rates are invariant, and so inferring infection generating rates in this way maintains consistency with cited sources.

The final two parameters, λ_h and λ_w, summarize the effect of lockdown policies on disease transmission and labor productivity, respectively. We choose λ_h to match the trajectory of cumulative cases in the U.S. under lockdown measures. As many recent studies have documented that case counts in randomized public testing for antibodies generally exceed reported cases by substantial multiples, we convert data on confirmed cases to actual cases by rescaling the reported data by a factor of 20. There is no consensus yet as to this value, but our choice of 20 non-reported cases for every reported infection is well within the range estimated by (Hortaçsu, Liu, and Schwieg, 2020). The productivity loss from lockdown policies, λ_w is calibrated to match the 32 percent decline in hours worked during the U.S. lockdown, as documented in the weekly labor market surveys of Bick and Blandin (2020).

We choose not to differentially calibrate the severity of lockdown policies across developing

---

[2]In particular, we take the highest cross-country peak infection rates as our target. The logic is that these countries likely have the least effective aggregate mitigation policies beyond precautions taken at the individual levels, and hence peak infection rates best reflect the behaviorally adjusted transmission rate in our model.

Var	Description	Value	Source or Target
η	Effect of infection on productivity	0.8	Asymptomatic cases
κ	Impact of hospital overuse on fatality	2	Glover et al. (2020)
λ_w	Effect of lockdown on productivity	0.68	Blandin and Bick (2020)
λ_h	Effect of lockdown on infection rate	0.75	U.S. cumulative infections
$\pi_y^{\mathbb{C}}$	Rate of young entering $\mathbb{C}$ from $\mathbb{I}$	3.43%	Ferguson et al. (2020)
$\pi_o^{\mathbb{C}}$	Rate of old entering $\mathbb{C}$ from $\mathbb{I}$	19.88%	Ferguson et al. (2020)
$\pi_y^{\mathbb{D}}$	Rate of young entering $\mathbb{D}$ from $\mathbb{C}$	2.76%	Ferguson et al. (2020)
$\pi_o^{\mathbb{D}}$	Rate of old entering $\mathbb{D}$ from $\mathbb{C}$	10.86%	Ferguson et al. (2020)
$\beta^{\mathbb{I}}$	Behavior-adjusted infection generating rate	2.0	Peak Infection Rates

Table 2: Calibration of Epidemiological Parameters

and developed countries, instead allowing cross-country differences to emerge endogenously from agents' optimizing behavior. Furthermore, while conclusive evidence is not yet available, existing cross-country data suggests that instituted lockdown policies were broadly similar at the aggregate. For instance, Figure A.3 uses cross-country data from Google's *Community Mobility Report* to document the average change in residential and workplace mobility during government lockdowns by country's level of economic development. In all cases, lockdown policies lead to a substantial increase in time spent at home and decrease in time spent at work, and the magnitude of these changes far outweighs any changes across a country's level of development. Studies of specific developing countries yield similar results. For instance, Figure A.6 shows that a drop in cross-district mobility in Ghana – a proxy for traveling to work – fell by roughly 25 percent at the onset of lockdown policies. Similarly, Figures A.4 and A.5 use high frequency labor market surveys in Ghana during the same period to show declines in total hours worked of 20 to 25 percent during lockdowns, broadly similar to numbers documented in the United States. Taken altogether, these data suggest lockdown technologies that are largely similar across countries and motivate our parsimonious parameterization.

Finally, Table 3 summarizes parameters which vary across advanced and developing countries. Further evidence of such variation is provided in Section 2. Sectoral total factor productivities are chosen to match the extent of labor market informality across levels of development. The utility value of living, $\bar{u}$, is set to match the statistical value of a life, and comes from Glover et al. (2020), renormalized to the average consumption level in each country type. The financial wedge in capital markets for developing countries is taken from Donovan (2018), whose model matches the low savings rates among poor African households. We take the iceberg costs resulting from low fiscal capacity from the study of tax collection efficiency in Ghana by

Var	Description	Advanced Economies	Developing Economies	Source or Target
A	Formal sectors TFP	3.0	0.15	1% labor informality in US
$\bar{u}$	Flow value of being alive	$11.4\bar{c}^{US}$	$11.4\bar{c}^{DEV}$	Glover et al. (2020)
χ	Spread b/w borrowing and lending	0	0.66%	Donovan (2019)
τ	Marginal tax rate	0.25	0.15	Besley and Persson (2013)
Δ	Iceberg cost in tax collection	1	2.22	Dzansi et al. (2013)
$\bar{B}$	Lockdown emergency transfers	1%	0.1%	Lockdown transfer programs
ω	Share of young in population	73%	92%	2018 ACS / World Bank
Π	Int' aid / natural resources revenue	0	10% of GDP	World Bank
Θ	Hospital capacity per capita	0.00042	0.00011	Glover et al. (2020) / WHO

Table 3: Calibration of Parameters Varying between Advanced and Developing Economies

Dzansi et al. (2018). The tax rates for the advanced and developing countries are taken from Besley and Persson (2013). We normalize away the effect of financial wedges and fiscal iceberg costs in the calibration of advanced economies, setting them to zero and one, respectively. Comparative values of age demographics and exogenous government revenue in the form of aid and natural resources are taken from the World Bank. In particular, the fraction of young (those between 15 and 65) is 73 percent in advanced economies and 92 percent in developing ones. We exclude those below age 15 from the analysis.

For advanced economies, we calibrate the size of emergency transfers, $\bar{B}$, to reflect the benefits paid out during the lockdown in the United States. Renormalized, these programs totaled about 1 percent of annual U.S. GDP every two weeks. While there is more heterogeneity in the developing world, recent evidence suggests that stimulus programs in Africa are planned to be about one-tenth the size of U.S. programs, as a share of domestic GDP (Collier et al., 2020). Accordingly, we set the level of transfers to 0.1 percent of GDP in developing countries.

The final parameters to be set govern the ICU hospital capacity in developing and developed countries. One challenge is that while many countries report hospital bed capacity, few developing countries distinguish explicitly between general hospital capacity and ICU capacity in the data. To address this, we assume the ratio of hospital beds to ICU beds is constant across countries, and calibrate Θ by adjusting WHO data on the availability of hospital beds in the top and bottom quartiles of country income levels (as in Figure 4) by the ratio of hospital beds to ICU beds taken from Glover et al. (2020).

4.2. Model Validation

Before reporting results, we check to ensure that our calibration strategy provides a reasonable fit to relevant moments in the data. We focus specifically on four salient moments crucial to the credibility of subsequent quantitative exercises. These include (1) the relative income levels of advanced and developing countries, (2) the relative size of the informal sector, (3) the fraction of hand-to-mouth consumers, and (4) the epidemiological dynamics of the pandemic.

Overall, in spite of its simplicity, the model does reasonably well at matching key non-targeted moments in low-income countries. It predicts that the fraction of workers employed in the informal sector is 68 percent in the developing economy (rather than 1 percent in the advanced economy) which is consistent with the evidence from Figure 3. The model's predicted fraction of hand-to-mouth households is 21 percent in the advanced economy and 69 percent in the developing economy. Direct analogues from the data are not readily available for the developing economy, though the advanced economy value is in line with the estimates of Kaplan, Violante, and Weidner (2014). The level of GDP per capita in the developing economy is \$1,100 per year, which is consistent with the average of the bottom quartile of countries.

Figure 6 provides some validation for the epidemiological dynamics of the model by comparing the progression of infections in the model to data currently available in the United States. In particular we plot the model's predicted cumulative infection rate and the number of infections in the data, assuming that there are 20 non-reported infections for every reported infection. This multiple for non-reported infections is within the range estimated by Hortaçsu et al. (2020). While real world data on the entire progression of infection counts is limited by questions on the prevalence of asymptomatic cases and the fact that the pandemic is still in its early days, our model's predictions are broadly consistent with currently available data. In future work, we plan to update these parameters as better measurements become available with time.

4.3. Lockdowns in Advanced and Developing Countries

Table 4 summarizes results on welfare, GDP, and fatalities under various lockdown policies in advanced and developed countries. The welfare and GDP entries report the percent change in each outcome variable relative to the country's pre-pandemic steady state levels, and fatalities are reported per hundred thousand people. In our baseline results, we consider aggregate policies which range in length from "no lockdown" to a 70-week lockdown, which lasts through most of the epidemic in the model. We set the maximum duration of the pandemic to 500 days, which is what previous studies have assumed about the amount of time it would take to develop a commercial vaccine. This is, of course, little more than a guess.

Figure 6: Model Fit of Cumulative Infection Cases in the United States
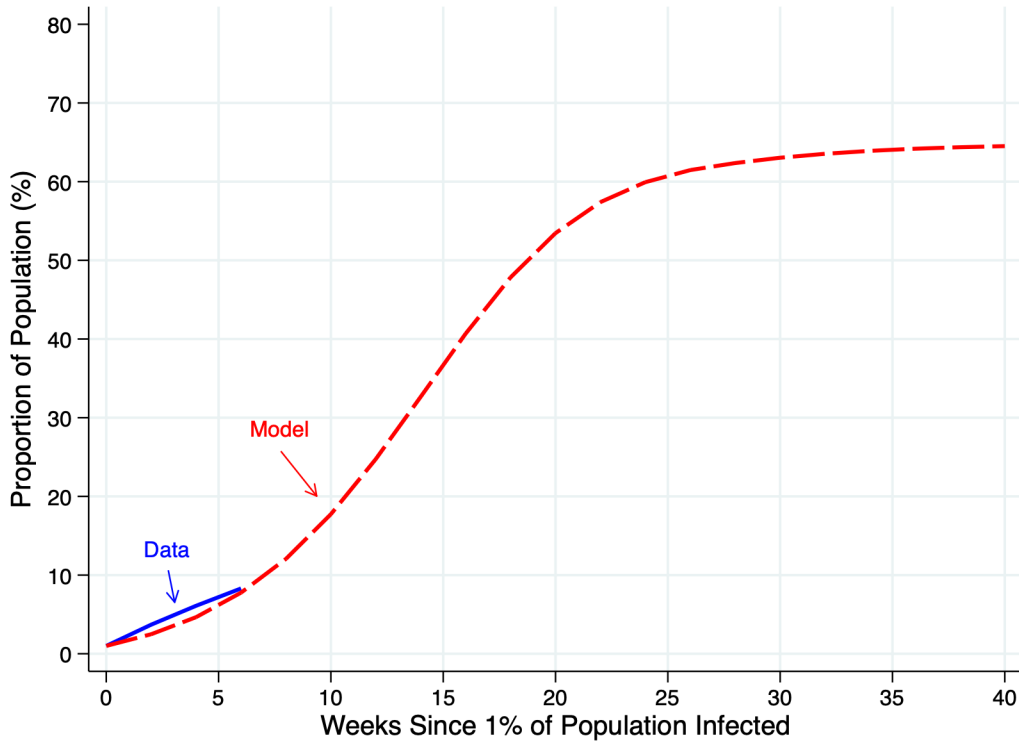
Note: Data on reported infections in the United States come from the European Centre for Disease Prevention and Control. We assume that there are 20 non-reported infections for every reported infection, as explained in the text.

As a benchmark, it is useful to consider what unfolds in each country when no aggregate lockdown policy is put in place. In advanced economies, doing nothing in response to the pandemic results in 1,102 deaths per hundred thousand people, a 1.8 percent contraction in GDP, and an 8.3 percent decline in aggregate welfare. The consequences of doing nothing in the developing world are about half as severe as in advanced countries. With no lockdown in the developing world, the pandemic leads to 412 deaths per hundred thousand people, a 1.1 percent contraction in GDP, and a 4.1 percent reduction in aggregate welfare. The lower cost for developing countries stems largely from their younger population, which is less likely to lose productivity from being sick or to die. Of course, in reality GDP losses may be even greater than what our model predicts due to features we have omitted, such as disruptions in supply chains (Bonadio, Huo, Levchenko, and Pandalai-Nayar, 2020), Keynesian demand channels (Guerrieri et al., 2020), input-output linkages (Baqaee and Farhi, 2020), or other forces.

While developing countries fare better than advanced ones in the absence of an aggregate policy response, they appear less able to effectively mitigate these negative outcomes through

Table 4: Predicted Effects of the COVID-19 Pandemic

	Lifetime Welfare (%)	GDP (%)	Fatalities per 100,000 People
Panel A: Advanced Economies			
No Lockdown	−8.3	−1.8	1,102
Twelve-Week Lockdown	−7.8	−8.9	1,026
Twenty-Eight-Week Lockdown	−5.5	−18.2	778
Seventy-Week Lockdown	−5.8	−32.8	767
Panel B: Developing Economies			
No Lockdown	−4.1	−1.1	412
Twelve-Week Lockdown	−4.0	−4.0	383
Twenty-Eight-Week Lockdown	−3.6	−8.2	340
Seventy-Week Lockdown	−3.9	−12.7	340

Note: This table reports changes in lifetime consumption equivalent welfare, GDP and fatalities per 100,000 people over the course of the COVID-19 pandemic relative to the pre-pandemic period. In all three lockdown scenarios, the length of the lockdown refers to the number of weeks since one percent of the population is infected. Advanced economies refers to the model's predictions for an economy calibrated to match features of countries in the top quartile of the world income distribution. Developing Economies refers to the model's predictions for an economy calibrated to the bottom quartile of the world income distribution.

lockdown policies. Specifically, while lockdowns always save lives, the costs are higher, and benefits more modest, in developing countries. For example, under the twenty-eight week lockdown, fatalities decline by around 30 percent in advanced economies, from 1,102 to 778 deaths per hundred thousand people. In developing economies, fatalities fall by only 18 percent, from 412 to 340 per hundred thousand people. The GDP losses under lockdown policies are greater in advanced economies than in developing ones, but proportionally less than differences in fatalities. Consequently, lockdown policies are about half as effective in developing countries as in advanced ones, saving half as many lives per percentage point of GDP lost. These asymmetries are reflected in the differential welfare benefits of lockdowns in the two economies, with moderate lockdowns reducing welfare losses by 35.9 percent in advanced economies, but only 13.8 percent in developing ones.
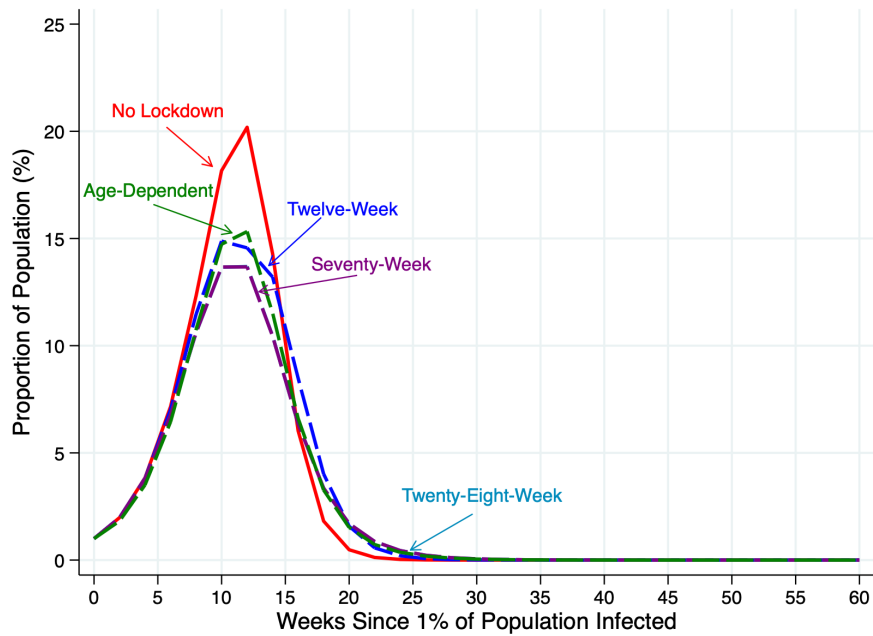
The lower relative efficacy of lockdowns in developing countries is true across the various

Figure 7: Simulated COVID-19 Infection Rates

(a) Advanced Economy
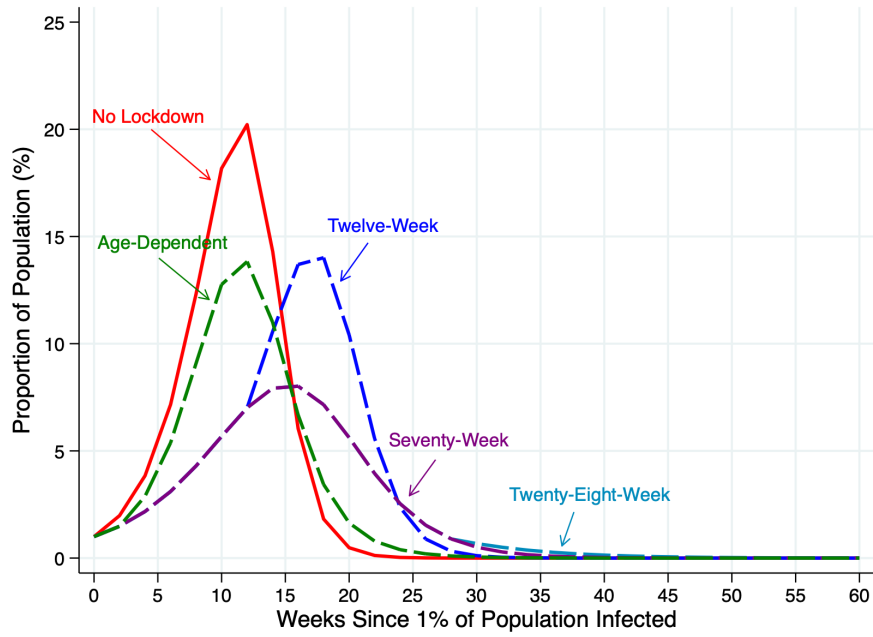
(b) Developing Economy

policy durations we consider. The main reason is that in the absence of robust transfer programs and in the presence of widespread informality, lockdowns do little to stymy the spread of infections in developing countries. Figure 7 shows this explicitly by plotting the trajectory of infections (including non-reported infections) under each policy for advanced and developing countries. At every duration, lockdowns yield less public health benefits in developing countries since they are widely flouted by low income individuals who move to the informal sector to offset earnings losses. Due to public health externalities in the spread of disease, a large non-complying population will erode the public health benefits to the economy at large. In developing countries, where TFP differences between the formal and informal sectors are lower and where tax-and-transfer programs are small and inefficient, the incentives to disregard lockdowns and move to the informal sector are substantial.

An important implication of this is that developing countries may have to worry about extending lockdown policies for too long. In terms of welfare, our model predicts that a 28-week lockdown is greatly preferable to a 70-week lockdown. However, for the advanced economy, the 70-week lockdown achieves welfare gains very similar to the 28-week lockdown. In particular, the gains from locking down for 70 weeks are 88 percent of the gains of locking down for 28 weeks. In contrast, the developing country loses much of the benefits if it locks down for too long. The 70-week gains are only around one third of the 28-week gains. This result suggests that developing countries may have to be more conservative than advanced ones when considering the length of their lockdowns.

4.4. Counterfactual Accounting

To understand the economics driving differential outcomes in advanced and developing countries, Figure 8 reports welfare, GDP, and fatality outcomes under a 28-week lockdown in the advanced economy as we sequentially endow it with the salient characteristics of developing countries identified in Section 2. The top panel reports counterfactual results for each channel in isolation. Since the magnitude of these individual channels implies substantial equilibrium interactions, the bottom panel adds each mechanism sequentially and cumulates their economic impact. For completeness, Figure A.7 reports the effect of each channel in the absence of any lockdown.

Three main lessons can be drawn from the counterfactual results. First, differences in mortality rates are driven overwhelmingly by the age distributions of advanced and developing countries. Endowing advanced economies with the age distribution of developing ones cuts the number of deaths by 65 percent on its own, accounting for more than the total difference in mortality between the two countries in our benchmark results. Taken individually, changing

the age distribution alone can account for nearly all of the differences between advanced and developing countries in the "no lockdown" scenario where aggregate policy does not respond to the pandemic. Furthermore, while the limited ICU capacity of developing countries increases mortality, it does so only modestly, with widespread labor market informality contributing even more to differences in death rates. The reason is that the effect of hospital congestion on survival probabilities is modest relative to the negative public health externalities in infection rates generated by workers flouting lockdown by moving into the informal sector. Interestingly, even though on its own the informality channel leads to more deaths than limited ICU capacity, the latter generates larger welfare losses since the increase in deaths are not accompanied by lower losses to GDP, as they are in the case of informality.

Turning to GDP outcomes, the most salient channel appears to be the extensiveness of labor market informality. On its own, endowing advanced economies with the labor market informality of developing ones serves to cut the output losses associated with lockdowns roughly in half. This is in part mechanical as the informal sector is not subject to lockdown measures, and so a larger informal sector means fewer workers being initially subject to the output losses associated with lockdowns. However, the informal channel also has important equilibrium interactions with the other mechanisms. Intuitively, the workers who find it optimal to circumvent lockdown orders by moving to the informal sector are those for whom the economic losses associated with lockdowns are greater than the health risks of the pandemic. Informality therefore allows these workers to self-select out of lockdown measures, incurring a smaller loss in income but greater health risk. Comparing the individual versus cumulative results, we see that this self-selection moderates the welfare and output losses associated with the other channels as well. This is especially true when one considers the interaction with ICU capacity and the share of hand-to-mouth; on their own, both lead to substantial welfare losses, but coupled with a reasonable option to move into informality, the losses are substantially lower. This result is far from obvious, as public health externalities in infection rates could have pushed outcomes in the opposite direction.

Given our heterogeneous agent setup, welfare outcomes under each scenario are driven not only by aggregate outcomes in GDP and fatalities, but also by how those changes are distributed in the population. For instance, on its own, lower fiscal capacity substantially increases the welfare costs of lockdowns even though it has negligible effects on output or fatalities. The reason of course is distributional, as low fiscal capacity limits the effectiveness of redistribution through tax-and-transfer and emergency relief programs. In this regard, the fiscal channel, labor market informality, and the age distribution play crucial roles in shaping welfare outcomes in the counterfactuals. The age distribution plays an important role in mitigating the effect of

Figure 8: Counterfactual Economies under 28-Week Lockdown

(a) Individual Contributions

| Welfare Loss (%) | GDP Loss (%) | Fatalities per 1,000 |

(b) Cumulative Contributions

| Welfare Loss (%) | GDP Loss (%) | Fatalities per 1,000 |

weaker ICU capacity in developing countries by mechanically shrinking the share of the vulnerable population. The fiscal and informality channels are important for blunting the welfare losses associated with lower incomes and a higher shares of hand-to-mouth consumers in developing countries. The informal channel appears more pronounced in this regard partially because, for simplicity, we have ruled out more sophisticated reoptimization by governments in shaping their tax-and-transfer programs in response to the pandemic's evolution.

In summary, our counterfactual analysis finds meaningful roles for each of the channels in shaping at least one of the outcomes on welfare, GDP, or mortality. Comparing the top and bottom panels of Figure 8 indicates that there are substantial equilibrium interactions between the mechanisms we study. This observation motivates the necessity for modeling-based approaches to studying the consequences of various COVID-19 policies. Our current analysis suggests the most important of these interactions in a developing world context are those between lower fiscal capacities, extensive labor market informality, and age demographics.

4.5. Heterogenous Effects of Lockdown

There is a growing body of research documenting substantial heterogeneity in who bears the costs and benefits associated with the pandemic's spread and the policy responses aimed at tackling it (e.g. Glover et al., 2020; Alon et al., 2020; Mongey et al., 2020). Accounting for such heterogeneity is important for understanding the economic incentives shaping behavioral responses to the pandemic and in evaluating the welfare consequences of different policy responses. The model accounts for these differences by including heterogeneity in age and incomes, both of which play an important role in propagating the economic consequences of the pandemic. Figure 9 summarizes the implications of this heterogeneity by plotting welfare changes by age and permanent productivity level (a proxy for permanent income in the model) under the lockdown and no-lockdown scenarios. The left column reports results for advanced economies, the right displays outcomes for developing ones.

Two trends are immediately clear from the Figure 9. First, in both countries, the welfare costs of the pandemic and welfare gains of lockdown policies accrue overwhelmingly to the old population. Second, heterogenous effects across the income distribution are relevant but less pronounced in both countries. Third, the welfare benefits of lockdowns broadly decrease with permanent productivity in advanced economies, but increase with income in developing ones. This final observation has an especially important implication for the aggregate effect of lockdowns in advanced and developing countries: in advanced countries, lockdown policies are generally progressive and benefit the poor more than the rich; in developing countries, lockdown policies appear to be regressive, and benefit the rich more than the poor.

Figure 9: Heterogeneous Effects of Lockdown (28-Weeks)

(a) Young, Advanced Economy

(b) Young, Developing Economy

(c) Old, Advanced Economy

(d) Old, Developing Economy
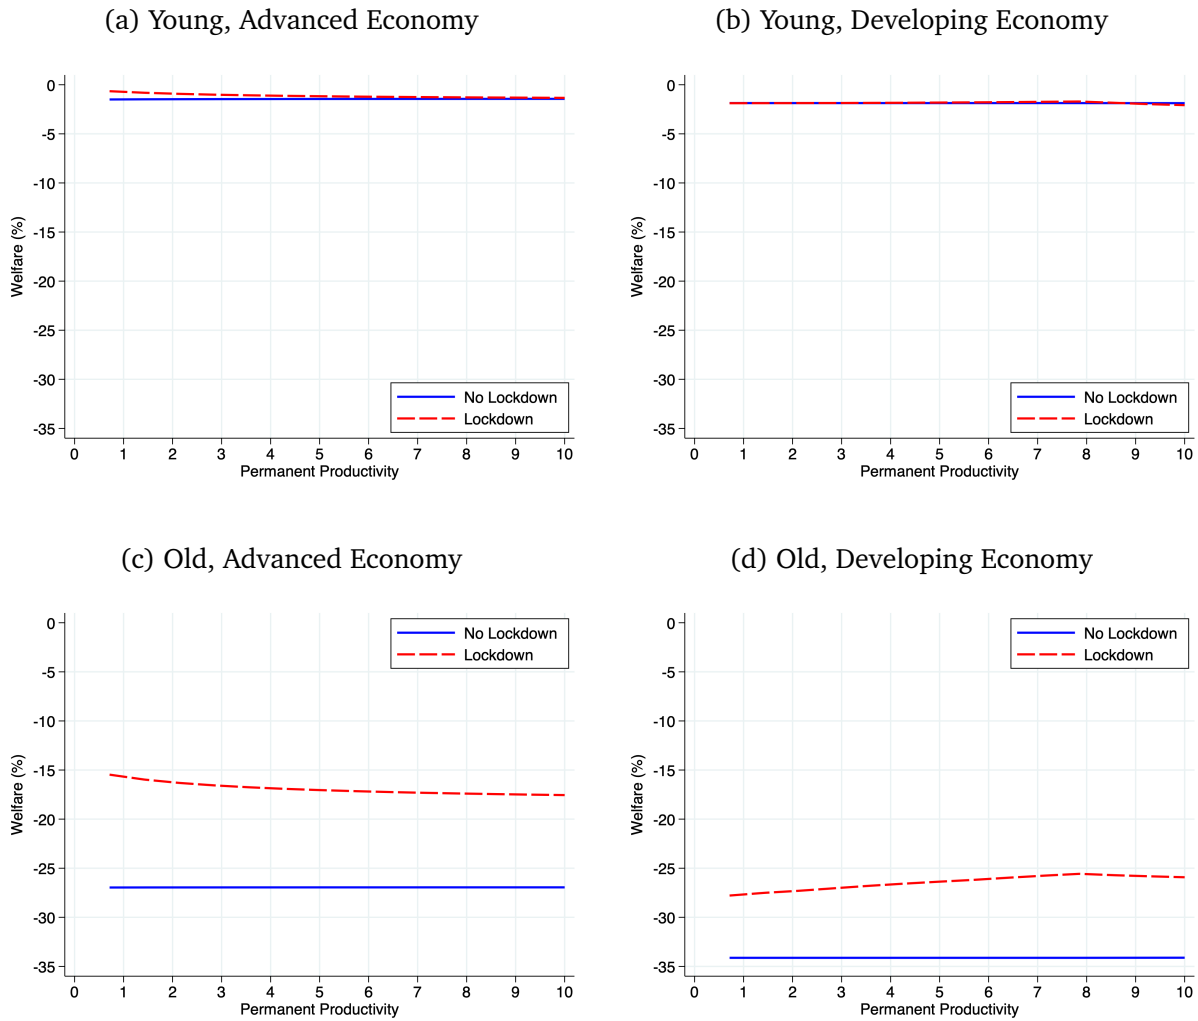
The fact that older individuals benefit disproportionately from lockdown measures is hardly surprising and is an immediate consequence of their substantially higher susceptibility to disease. On average, lockdown policies reduce welfare losses amongst the old in both countries by about ten percentage points. The young also gain from the public health benefits of lockdowns, though meaningfully less due to their lower inherent susceptibility, and these gains appear largely offset by the economic losses accompanying the lockdown.

A more nuanced result is that in advanced economies the welfare benefits of a lockdown falls with income levels, while in developing countries they broadly increase. This outcome is a consequence of the subtle interactions between fiscal capacity and informality with the differing income levels in the two countries. When the lockdown goes into effect, governments substantially increase transfers which implicitly transfer welfare from rich households to poor

ones in both countries. As rich countries implement larger programs and have higher fiscal capacity, the overall impact of this redistribution is substantial, giving rise to the negative slope in welfare gains across the income distribution. In developing countries, these emergency transfer programs are smaller and waste substantial resources because of the lower fiscal capacity of governments in the developing world. Widespread labor market informality in developing countries compounds these effects, as many low income workers seek to offset their income losses by moving into the informal sector. While partly alleviating economic hardships, the move to informality increases the exposure of poor households to disease, offsetting the beneficial health effects of the lockdown relative to their richer compatriots who have sufficient assets to avoid informality. On balance, the informality channel supersedes the redistributive effects of transfers in developing countries, giving rise to the upward sloping curve in the figure.

4.6. Effectiveness of Age-Dependent Policies

The large and heterogenous effects of lockdowns suggest that targeted policies may be more effective than the benchmark lockdowns studied above. Specifically, while blanket lockdowns impose costs on the entire economy, the benefits accrue overwhelmingly to the old. Age-dependent lockdown policies which focus on shielding only the old could therefore deliver similar benefits as blanket lockdowns, but at a much lower cost. Several other recent research papers have similarly argued for the potential advantages of such age-targeted programs in advanced economies (Acemoglu et al., 2020; Bairoliya and Imrohoroglu, 2020). Our analysis suggests that these targeted policies may be even more potent in the developing world, which is more sensitive to the economic costs of lockdowns and where the old constitute a smaller and more vulnerable share of the population.

To evaluate the efficacy of more targeted programs in the developing world, we analyze how outcomes change when lockdown policies are targeted exclusively at the old. To facilitate comparability, we assume that the lockdown technology parameters are the same, but do not apply to the young population, and keep fixed the total amount of money spent on emergency transfers, B_t, simply redistributing the excess resources exclusively to the old population.

In both advanced and developing countries, age-dependent policies appear to greatly increase the efficacy of lockdown programs when assessed using lives saved per percentage of lost GDP. Broadly, the added potency of age-dependent policies appears greater in short lockdowns than in long ones, and far greater in developing countries than advanced ones. The added effectiveness of age-targeting for short lockdowns stems in part from the fact that short-lockdowns generally do little to stymy the spread of disease but still lead to large economic costs. As a result, introducing age targeting in these shorter programs proportionately increases the lives

Table 5: Lives Saved per 100,000 People per Unit of GDP Lost

	Advanced Economy		Developing Economy	
	Benchmark	Age-dependent	Benchmark	Age-dependent
Twelve-Week	10.6	54.0	9.8	148.3
Twenty-Eight-Week	19.8	54.0	10.2	95.2
Seventy-Week	10.8	29.8	6.2	44.5

Note: This table reports the number of lives saved per 100,000 people per unit of GDP lost to the lockdown. The Benchmark columns refers to the effects of a blanket lockdown, and the Age-Dependent columns refer to lockdowns which keep only the older population under lockdown.

saved by a greater amount than in longer lockdowns, which are already effective on this dimension.

More germane to our purposes is the fact that age-dependent policies appear to be far more potent in increasing the efficacy of lockdowns in the developing economies than in advanced ones. In advanced economies, a 28-week lockdown with age targeting saves 54 lives per hundred thousand people for every unit of lost GDP. In developing economies, the same ratio is 95.2. Consequently, age-targeting in advanced economies increases lives saved per percentage of GDP lost by a factor of 3 to 5 over non-targeted policies, while in developing countries, increases are on the order of 7 to 15. The result stems in part from the fact that developing countries have far smaller old populations than advanced ones, so targeted policies have the added benefit of partially offsetting the weaker fiscal capacity of developing countries by concentrating transfers on a smaller population. Furthermore, targeted policies are better at keeping workers in the formal sector in developing countries, muting output losses and boosting government revenue relative to blanket lockdowns. Taken together then, we conclude that targeted policies are more potent in the developing world because in equilibrium they serve to mitigate the negative drag from other channels that are otherwise reinforced in blanket lockdown policies.

5. Conclusion

This paper provides a preliminary quantitative analysis of how lockdown policy should differ between developing and developed economies. Developing economies have different charac-

teristics that suggest differing lockdown policies from the west, including younger populations, larger informal sectors and lower healthcare and fiscal capacity. Our quantitative macroeconomic model predicts that blanket lockdowns are generally less effective in developing countries, and save fewer lives per unit of lost economic output. Nevertheless, in our simulations, blanket lockdowns still prove more effective in lowering the welfare costs of the pandemic than having no lockdown at all. The most effective type of lockdown for developing countries, according to our analysis, is one that locks down only the older population, sending transfers only to them. These age-dependent lockdowns have potent effects in our model, saving more lives per percent of GDP lost than the same policies in richer countries.

Our quantitative results for welfare, GDP losses, and fatality rates from the pandemic are best viewed as preliminary given that the pandemic is still, unfortunately, in its early stages, and reliable data are still scarce. Our conclusions about the differential effects of COVD-19 between advanced and developing economy may prove more enduring. There can be little doubt that developing economies have vastly younger populations, much larger informal economies and less fiscal capacity than advanced economies. These basic differences in demographic and economic structure point to optimal lockdown policies that are age-dependent, focusing on keeping just the older population under lockdown, and letting others resume normal economic activity, to the extent that it is possible.

References

ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON (2020): "A Multi-Risk SIR Model with Optimally Targeted Lockdown," NBER Working Paper No. 27102.

AIYAGARI, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving*," *The Quarterly Journal of Economics*, 109, 659–684.

AKTAY, A., S. BAVADEKAR, G. COSSOUL, J. DAVIS, D. DESFONTAINES, A. FABRIKANT, E. GABRILOVICH, K. GADEPALLI, B. GIPSON, M. GUEVARA, ET AL. (2020): "Google COVID-19 Community Mobility Reports: Anonymization Process Description (version 1.0)," *arXiv preprint arXiv:2004.04145*.

ALON, T. M., M. DOEPKE, J. OLMSTEAD-RUMSEY, AND M. TERTILT (2020): "The Impact of COVID-19 on Gender Equality," NBER Working Paper No. 26947.

ALVAREZ, F. E., D. ARGENTE, AND F. LIPPI (2020): "A Simple Planning Problem for COVID-19 Lockdown," NBER Working Paper No. 2698.

BAIROLIYA, N. AND A. IMROHOROGLU (2020): "Macroeconomic Consequences of Stay-at-Home Policies During the COVID-19 Pandemic," *Covid Economics*, 13, 71–90.

BAKER, S. R., N. BLOOM, S. J. DAVIS, AND S. J. TERRY (2020): "COVID-Induced Economic Uncertainty," NBER Working Paper No.26983.

BAQAEE, D. R. AND E. FARHI (2020): "Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis," Unpublished Working Paper, Harvard University.

BERGER, D. W., K. F. HERKENHOFF, AND S. MONGEY (2020): "An SEIR Infectious Disease Model with Testing and Conditional Quarantine," NBER Working Paper No. 2690.

BESLEY, T. AND T. PERSSON (2009): "The Origins of State Capacity: Property Rights, Taxation, and Politics," *American Economic Review*, 99, 1218–44.

——— (2013): "Taxation and Development," in *Handbook of Public Economics*, Elsevier, vol. 5, 51–110.

BEWLEY, T. (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16, 252 – 292.

BICK, A. AND A. BLANDIN (2020): "Real Time Labor Market Estimates During the 2020 Coronavirus Outbreak," Unpublished Manuscript, Arizona State University.

BONADIO, B., Z. HUO, A. A. LEVCHENKO, AND N. PANDALAI-NAYAR (2020): "Global Supply Chains in the Pandemic," Unpublished Working Paper, University of Texas at Austin.

CASELLI, F. AND W. J. COLEMAN (2006): "The World Technology Frontier," *American Economic Review*, 96, 499–522.

ÇAKMAKH, C., S. DEMIRALP, Ş. KALEMLI-ÖZCAN, S. YESILTAS, AND M. A. YILDIRIM (2020): "COVID-19 and Emerging Markets: An Epidemiological Multi-Sector Model for a Small Open Economy with an Application to COVID-19," NBER Working Paper No. 27191.

CHANG, R. AND A. VELASCO (2020): "Economic Policy Incentives to Preserve Lives and Livelihoods," *Covid Economics*, 14, 33–56.

COLLIER, P., E. J. SIRLEAF, AND N. OKONJO-IWEALA (2020): "Global Leadership to Support Africa's Response to COVID-19," https://www.theigc.org/event/global-leadership-to-support-africas-response-to-covid-19/, accessed 17/05/20.

DONOVAN, K. (2018): "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences," Unpublished Working Paper, Yale University.

DZANSI, J., A. JENSEN, D. LAGAKOS, I. OTOO, H. TELLI, AND C. ZINDAM (2018): "Survey of Local Government Revenue Mobilization Capacity in Ghana, 2017: Summary and Policy Implications," International Growth Centre Project Report S-33417-GHA-2.

EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT (2020): "The Macroeconomics of Epidemics," NBER Working Paper No. 26882.

FEENSTRA, R. C., R. INKLAAR, AND M. P. TIMMER (2015): "The Next Generation of the Penn World Table," *American Economic Review*, 105, 3150–82.

FENG, Y., D. LAGAKOS, AND J. E. RAUCH (2018): "Unemployment and Development," NBER Working Paper No. 25171.

FERGUSON, N. M., D. LAYDON, AND G. NEDJATI GILANI (2020): "Report 9: Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand," Unpublished Working Paper, Imperial College London.

FLODEN, M. AND J. LINDÉ (2001): "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?" *Review of Economic Dynamics*, 4, 406–437.

FLOWMINDER (2020): "Mobility Analysis to Support the Government of Ghana in Responding to the COVID-19 Outbreak," Https://web.flowminder.org/publications/mobility-analysis-to-support-the-government-of-ghana-in-responding-to-the-covid-19-outbreak.

GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RÍOS-RULL (2020): "Health Versus Wealth: On the Distributional Effects of Controlling a Pandemic," NBER Working Paper No. 27046.

GOLLIN, D. (2002): "Getting Income Shares Right," *Journal of Political Economy*, 110, 458–474.

GREENWOOD, J., P. KIRCHER, C. SANTOS, AND M. TERTILT (2019): "An Equilibrium Model of the African HIV/AIDS epidemic," *Econometrica*, 87, 1081–1113.

GUERRIERI, V., G. LORENZONI, L. STRAUB, AND I. WERNING (2020): "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" NBER Working Paper No. 26918.

HORTAÇSU, A., J. LIU, AND T. SCHWIEG (2020): "Estimating the Fraction of Unreported Infections in Epidemics with a Known Epicenter: an Application to COVID-19," NBER Working Paper No. 27028.

HUGGETT, M. (1993): "The Risk-Free Rate in Heterogeneous-agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969.

ICTD/UNU-WIDER (2019): "Government Revenue Dataset," https://www.wider.unu.edu/project/government-revenue-dataset, Last accessed on 2020-05-14.

JENSEN, A. (2019): "Employment Structure and the Rise of the Modern Tax System," NBER Working Paper No. 25502.

JONES, C. I. AND P. J. KLENOW (2016): "Beyond GDP? Welfare across Countries and Time," *American Economic Review*, 106, 2426–57.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2020): "Pandemics According to HANK," Unpublished Presentation, LSE.

KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, spring.

KRUEGER, D., H. UHLIG, AND T. XIE (2020): "Macroeconomic Dynamics and Reallocation in an Epidemic," *Covid Economics*, 16.

LAGAKOS, D., A. M. MOBARAK, AND M. E. WAUGH (2019): "The Welfare Effects of Encouraging Rural-Urban Migration," Unpublished Working Paper, Yale University.

LAGAKOS, D. AND M. E. WAUGH (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 103, 948–80.

MALMBERG, H. (2018): "How Does the Efficiency of Skilled Labor Vary across Rich and Poor Countries? An Analysis Using Trade and Industry Data," Unpublished Manuscript, University of Minnesota.

MONGEY, S., L. PILOSSOPH, AND A. WEINBERG (2020): "Which Workers Bear the Burden of Social Distancing Policies: New Evidence," Becker Friedman Institute Working Paper 2020-51.

PRIN, M. AND H. WUNSCH (2012): "International Comparisons of Intensive Care: Informing Outcomes and Improving Standards," *Current Opinion in Critical Care*, 18, 700.

ROY, A. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–146.

TODA, A. A. (2020): "Susceptible-Infected-Recovered (SIR) Dynamics of COVID-19 and Economic Impact," *arXiv:2003.11221*.

ULYSSEA, G. (2018): "Firms, Informality, and Development: Theory and Evidence from Brazil," *American Economic Review*, 108, 2015–47.

ZHAO, B., K. STORESLETTEN, AND F. ZILIBOTTI (2019): "Business Cycles During Structural Change: Arthur Lewis? Theory from a Neoclassical Perspective," Unpublished Working Paper, Yale University.

Appendix

A. Appendix Figures

Figure A.1: Fraction of the Population Older than Age 65

Note: This figure plots the proportion of population ages over 65 and above as a percentage of total population across 162 countries. GDP per capita is from Penn World Table 9.1 (Feenstra et al., 2015). Population data is World Bank staff estimates using the World Bank's total population and age/sex distributions of the United Nations Population Division's World Population Prospects: 2019 Revision.

Figure A.2: Share of Self-Employed Workforce

Note: This figure plots the self-employed workers as a percentage of total employment across 153 countries. Self-employed workers are those workers who, working on their own account or with one or a few partners or in cooperative, hold the type of jobs defined as a "self-employment jobs." i.e. jobs where the remuneration is directly dependent upon the profits derived from the goods and services produced. Self-employed workers include four sub-categories of employers, own-account workers, members of producers' cooperatives, and contributing family workers. GDP per capita is from Penn World Table 9.1 (Feenstra et al., 2015). Self-employment data is from ILOSTAT database.

Figure A.3: Changes in Mobility Across Countries During Lockdown Periods


Note: This figure plots the average percentage changes of the mobility metric in the 'Places of Residence' and 'Workplace' categories in the Google Community Mobility Report (Aktay et al., 2020), during the lockdown periods for the 65 countries which had implemented or are implementing lockdown. GDP per capita is from Penn World Table 9.1 (Feenstra et al., 2015). The average across all 65 countries is 23.44 percent. The slope of the fitted line is 1.52, with p-value of 0.354 for the 'Workplace' category. For the 'Places of Residence' category, the slope of the fitted line is -1.52, with p-value of 0.083.

Figure A.4: Employment Rate in Ghana Around the Lockdown Period
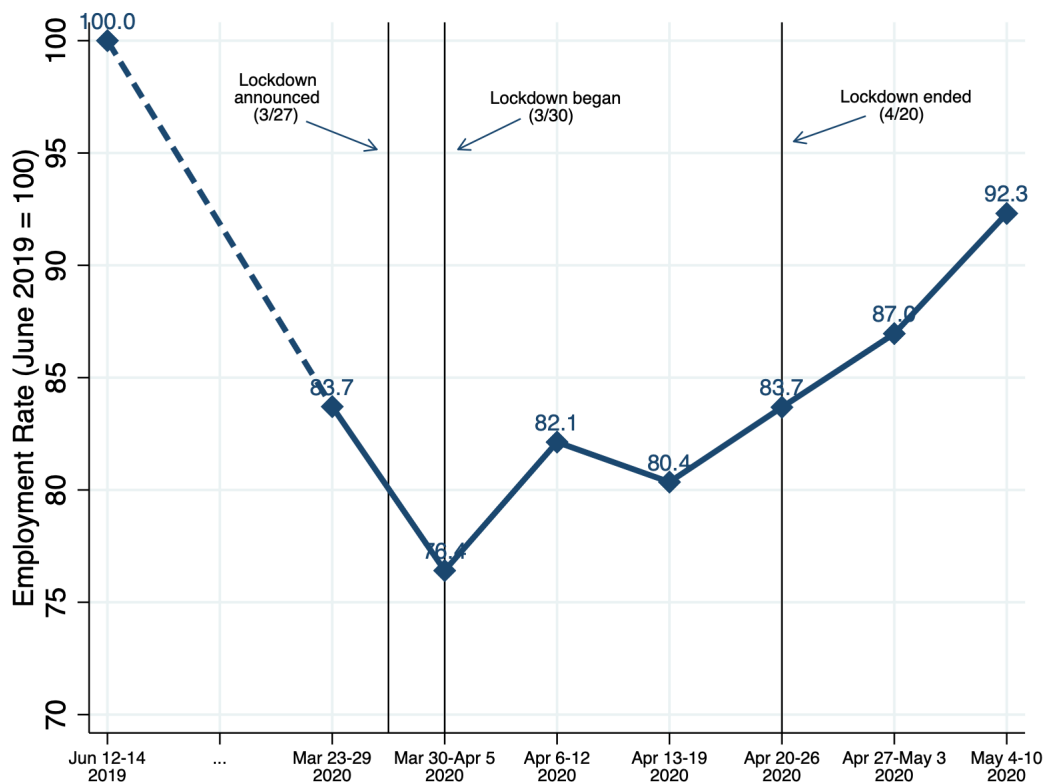
Note: This figure plots the employment in Accra and Kumasi, the two biggest cities in Ghana, proxied by Google survey data. The survey question was "In the last WEEK, how many hours did you work for pay or profit?". Employment rate is the proportion of people who reported they had worked in last week. The employment rate from the same survey conducted in June 2019 was normalized to one. Sample size was 266 in March 23-29, 342 for April 1-5, 281 for April 6-12, 302 for April 13-19, and 500 for the rest.

Figure A.5: Hours Worked in Ghana Around the Lockdown Period

Note: This figure plots the employment in Accra and Kumasi, the two biggest cities in Ghana, proxied by Google survey data. The survey question was "In the last WEEK, how many hours did you work for pay or profit?". Average hours per worker is calculated by taking the average of hours among those who reported they had worked in the last week. Average hours per adult is calculated by taking the average of hours including those who reported they had not worked in the last week. Sample size was 266 in March 23-29, 342 for April 1-5, 281 for April 6-12, 302 for April 13-19, and 500 for the rest.

Figure A.6: Mobility in Ghana Around the Lockdown Period
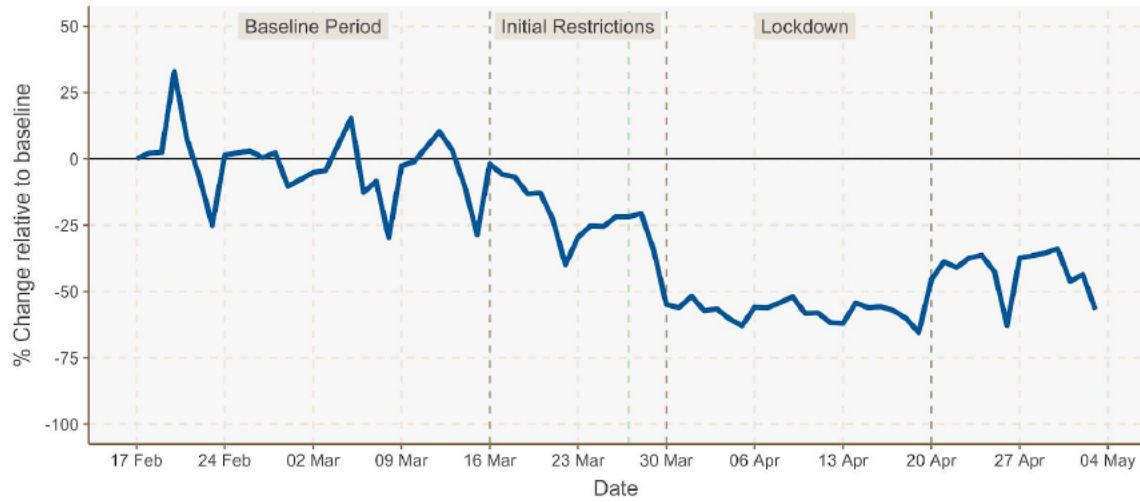
Note: This figure plots the percentage change in the number of trips between any two districts in Greater Accra, Ghana in each day, relative to the baseline value. The baseline value is calculated as the median value of the metric during the four weeks prior to the introduction of the first restriction on March 16th. The trips mainly comprise short-distance, routine daily trips that correspond to activities such as commuting to work, shopping, and entertainment. Anonymized and aggregated mobile phone data from Vodafone Ghana, analysis by Flowminder. Source: Flowminder (2020)

Figure A.7: Counterfactual Economies in No Lockdown Scenario

B. Appendix Tables

Table B.1: Counterfactual Exercise

	Lifetime Welfare (%)	GDP (%)	Fatalities per 100,000 People
Panel A: Advanced Economy at Baseline			
No Lockdown	−8.34	−1.76	1,102
Twenty-Eight-Week Lockdown	−5.45	−18.17	778
Panel B: Advanced economy with lower fiscal capacity			
No Lockdown	−8.30	−1.76	1,102
Twenty-Eight-Week Lockdown	−6.28	−18.17	778
Panel C: Advanced economy with large informal sector			
No Lockdown	−10.40	−1.76	1,102
Twenty-Eight-Week Lockdown	−7.32	−8.96	859
Panel D: Advanced economy with younger population			
No Lockdown	−3.11	−1.12	395
Twenty-Eight-Week Lockdown	−2.12	−17.81	271
Panel E: Advanced economy with lower hospital capacity			
No Lockdown	−8.44	−1.78	1,131
Twenty-Eight-Week Lockdown	−5.68	−18.19	825
Panel F: Advanced economy with more hand-to-mouth households			
No Lockdown	−8.28	−1.76	1,102
Twenty-Eight-Week Lockdown	−5.40	−18.17	778

Table B.2: ICU Bed Availability Across Countries

Country	ICU beds per 100,000 population	Per capita healthcare cost
United States	20.0-31.7	$7,164
Canada	13.5	$3,867
Denmark	6.7-8.9	$3,814
Australia	8.0-8.9	$3,365
South Africa	8.9	$843
Sweden	5.8-8.7	$3,622
Spain	8.2-9.7	$2,941
Japan	7.9	$2,817
UK	3.5-7.4	$3,222
New Zealand	4.8-5.5	$2,655
China	2.8-4.6	$265
Trinidad and Tobago	2.1	$1,237
Sri Lanka	1.6	$187
Zambia	0	$80

Source: Table 1 in Prin and Wunsch (2012). Healthcare cost includes all public and private expenditures, not limited to critical care.