Economic Activity during the COVID-19 Pandemic: A Model with "Acquired Immunity"

Juan Esteban Carranza, Juan David Martin, and Álvaro José Riascos

We calibrate a macroeconomic model with epidemiological restrictions using Colombian data. The key feature of our model is that a portion of the population is immune and cannot transmit the virus, which improves substantially the fit of the model to the observed contagion and economic activity data. The model implies that during 2020, government restrictions and the endogenous changes in individual behavior saved around 15,000 lives and decreased consumption by about 4.7 percent. The results suggest that most of this effect was the result of government policies. (JEL E1, I1, H0)

1 INTRODUCTION

In this article we formulate and calibrate a dynamic macroeconomic model in which optimizing agents respond to the risk of contagion and restrictions imposed by the government during the recent public health crisis. Our model is similar to the model by Eichenbaum, Rebelo, and Trabandt (2020; henceforth ERT), except for the inclusion of a modified epidemiological model that incorporates the possibility of exogenous immunity to contagion. We calibrate the model with Colombian data and use it to simulate counterfactual policies.

The original ERT model was the first of a wave of new articles using variations of a simple susceptible-infected-recovered (SIR) epidemiological model to account for the endogenous risk of contagion faced by economic agents. Other articles with similar approaches include Atkeson (2020); Alvarez, Argente, and Lippi (2020); Acemoglu et al. (2020); and Berger, Herkenhoff, and Mongey (2020). In these models, both contagion and economic activity are the results of a dynamic programming problem in which agents maximize their intertemporal utility, accounting for the risk of contagion over the course of an epidemic.

In the SIR model and its variations (based on seminal work by Kermack and McKendrick, 1927), an epidemic runs its course as it infects individuals who then become immune. The epidemic ends

when the population reaches "herd immunity," which occurs when enough people are immune and the virus cannot spread anymore. The main drawback of these models is the difficulty they have in fitting the observed COVID-19 contagion data. In particular, the standard model predicts very high numbers of both infections and deaths compared with the relatively low numbers observed in the data.

Our main contribution to understanding the pandemic is the modification of the SIR model to allow for the presence of individuals who are unaffected by the virus and who become immune over time at an exogenous rate. As our results show, the presence of this "immune" population helps to fit the model to the observed data. In particular, the model replicates well the rapid decline of observed deaths after the infection of a relatively low portion of the population during the first wave of the pandemic.

In contrast to ERT, we calibrate the model to match measures of both the pandemic and economic activity. We follow ERT in modeling government restrictions as a consumption tax, which induces consumers to cut back in their consumption activities, but we actually calibrate the parameterization of this tax.

The calibrated model predicts that almost all infections in Colombia will have already occurred by December 2020 and that the economy will be back on its long-term path by mid-2021. Our simulations suggest that government restrictions decreased yearly 2020 consumption by around 3 percent and saved around 10,000 lives. Without government restrictions, the economy would have still faced a 1 percent contraction, generated by consumers cutting back on consumption and labor to avoid contagion.

The model can be easily extended to accommodate successive contagion waves by allowing agents who become immune to become susceptible again. In the model, new waves can be triggered by making immunity disappear after an exogenous number of weeks or by letting immune agents become susceptible every period at an exogenous rate. This type of modeling would capture the possibility that, for example, either variations of the original virus appear or antibodies acquired from a mild infection disappear. Because the reasons individuals may or may not be immune are not well understood, we focus on the modeling of one wave of the pandemic and leave extensions for further research.

The article is organized as follows: Section 2 describes the model and its calibration. Section 3 contains the baseline results and counterfactual simulations. Section 4 concludes.

2 THE MACRO-SIOD MODEL

2.1 Description of the Model

The Model of the Pandemic with Endogenous Contagion. As we indicated above, we follow closely ERT and formulate a model with infinitely lived consumers who choose to allocate time into labor and consumption to maximize their lifetime expected utility. Their choices determine both the level of observed economic activity and the probability of contagion.

There are four types of agents in the model, depending on their exposure to the infection. These include susceptible (S), infected (I), survivor (O), and dead (D)—SIOD. A survivor agent can in turn be immune (M) or recovered (R). Notice that, in contrast to the standard SIR model embedded in ERT, we add the additional immune type (M), which is a type of patient that develops no symptoms and is not infectious.

In the model, agents become immune over time at an exogenous rate, which we calibrate.[1] Once they become immune, the immune agents behave similarly to recovered agents. The difference between an immune agent and a susceptible agent who becomes infected and then recovers is that the immune agent is never contagious and therefore never propagates the pandemic.

The existence of this type of "immune" agent is consistent with the increasing evidence of pre-existing immunity to the SARS-CoV-2 virus among a significant part of the population (Doshi, 2020). More broadly, our definition of immunity is consistent with other biological mechanisms that are not well understood yet, the details of which fall beyond the scope of this article. For example, this "immunity" is equivalent to situations in which individuals become infected but transmit the virus at variable rates (Adam, 2020). More specifically, if an individual does not transmit the virus after infection, they are "immune" according to our definition.

Denote T_t as the flow of newly infected individuals at time t, which depends on the probability that the stock of susceptible agents S_t becomes infected while interacting with the stock of infected agents I_t during consumption activities, work, or other activities, denoted π_1, π_2, and π_3, respectively; that is,

(1)
$$T_t = S_t\left(1-\pi_m\right)I_t\left(\pi_1 C_t^S C_t^I + \pi_2 N_t^S N_t^I + \pi_3\right),$$

where C_t^j and N_t^j correspond to the consumption and work hours of j-type agents with $j = S,I$.

The stock of infected agents evolves over time depending on (1) as follows:

(2)
$$I_{t+1} = I_t\left(1-\pi_{dt}-\pi_r\right)+T_t,$$

where π_{dt} and π_r are the probabilities of death and recovery, conditional on I. The probability of death changes over time depending on the capacity restrictions of the health care system, denoted as ξ in the following equation:

(3)
$$\pi_{dt} = \pi_d + 1_{\{I_t>\xi\}}\kappa I_t^2,$$

where the probability increases in a quadratic way when the number of infected individuals surpasses the capacity ξ, while κ is a parameter to be calibrated.

The main innovation of our work is the inclusion of type M agents who become immune over time without being infected and whose stock evolves, as follows:

(4)
$$M_{t+1} = M_t + \pi_m S_t,$$

where π_m is the exogenous probability of becoming immune. Notice that we assume that agents become immune over time at a constant rate, probably as a result of their exposure to the virus through interactions with infected agents. This simplifying assumption recognizes the fact that this immunity is not yet well understood. Notice that setting $M_0 = 0$ and $\pi_m = 0$ yields the standard SIR model, which we can also calibrate as a particular case in our model.[2]

To complete the description of the epidemiological model, the following equations describe the evolution of the stock of agent types D, R, and S:

(5)
$$D_{t+1} = D_t + \pi_{dt} I_t,$$
$$R_{t+1} = R_t + \pi_r I_t, \text{ and}$$
$$S_{t+1} = S_t (1 - \pi_m) - T_t,$$

where we assume that the initial stock of susceptible agents is the initial population $S_0 = Pop_0$ and that the initial stock of infected agents is nil; that is, $I_0 = \epsilon > 0$, a small portion of the population, which is calibrated.

The Economic Model. The economic problem of agents is the maximization of their lifetime utility through consumption and work decisions. Their choices determine the total number of hours devoted to consumption C_t and work N_t, which in turn determines the transition across agent types described above and the endogenous probability of contagion faced by a susceptible agent; that is,

(6)
$$\tau_t = \pi_1(c_t^S)(I_t C_t^I) + \pi_2(n_t^S)(I_t N_t^I) + \pi_3(I_t).$$

On the supply side, we assume there is a continuum of competitive, identical firms that use only labor and maximize period-by-period profits; that is,

(7)
$$\Pi_t = AN_t - w_t N_t,$$

where A is a productivity parameter and w_t is the competitive wage. In this closed economy, in equilibrium, total production must be equal to total consumption.

At each time t, agents are identified by their infection status $j \in S, I, O$. The problem of each agent type j at $t = t_0$ is given by

(8)
$$\max_{c,n} U_{t_0}^j = \sum_{t=t_0}^{\infty} \beta^{t-t_0} u(c_t^j, n_t^j),$$

where

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

and θ is a parameter to be calibrated.

The agent faces a budget constraint given by

(9)
$$c_t^j = \phi^j w_t n^j - \mu_t c_t^j + \Gamma_t,$$

where w_t is the hourly wage, μ_t is an exogenous consumption tax rate, and Γ_t is a government transfer that does not depend on the type of agent and that balances government finances. The parameter ϕ_t^j is a measure of the work ability of agent j due to infection, with the assumption that $\phi^S = \phi^O = 1$ and $\phi^I \leq 1$.

Thus, the probability of contagion τ_t affects the economic decisions of each type of agent through their expected lifetime utility at period t; that is,

(10)
$$U_t^S = u\left(c_t^S, n_t^S\right) + \beta\left[\left(1 - \pi_m\right)\left(\left(1 - \tau_t\right)U_{t+1}^S + \tau_t U_{t+1}^I\right) + \pi_m U_{t+1}^O\right],$$

$$U_t^I = u\left(c_t^I, n_t^I\right) + \beta\left[\left(1 - \pi_r - \pi_{dt}\right)U_{t+1}^I + \pi_r U_{t+1}^O\right], \text{ and}$$

$$U_t^O = u\left(c_t^O, n_t^O\right) + \beta U_{t+1}^O.$$

The tax rate μ_t plays an important role in the model because it absorbs all the restrictions imposed by the government. Policies such as quarantines and limits to the gathering of people are rationalized in the model as a consumption tax. In our application, we will identify this parameter using the period-by-period measure of economic activity.

In equilibrium, it must hold that all agents maximize their expected lifetime utility, firms maximize profits, the government balances its budget, and both labor and goods markets clear. We specify the equilibrium conditions in the appendix.

2.2 Calibration

We calibrate the model using weekly Colombian data matching the observed path of the COVID-19 pandemic throughout 2020. For any parametrization of the model, it is solved using a backward induction algorithm that involves (i) finding the optimal sequence of working hours for every type of agent along 250 weeks and then, similarly to ERT, (ii) computing the rest of the equilibrium sequences using the first-order conditions of the model.

We follow closely the criteria in ERT in choosing the parameters of the economic model, which we show in Table 1. We then calibrate the parameters $\pi = \{\pi_1, \pi_2, \pi_3, \pi_d, \pi_r, \pi_m\}$, which determine the dynamics of the pandemic. Moreover, we also calibrate the maximum percentage of the population that gets infected, after which there is "herd immunity."

As pointed out, another difference between our model and ERT is the treatment of the tax rate, which, as we pointed out, reflects the restrictions imposed by the government to contain the pandemic. We set $\mu_t = \mu$ for $t = 1$ until $t = 19$, which corresponds to August 31, 2020, when the national lockdown imposed by the government was officially ended. This policy had been in place since March 23 ($t = -4$). For $t \geq 20$, we set $\mu_t = 0.9\mu_{t-1}$ so that restrictions decreased gradually toward zero.

As indicated above, we need to calibrate the parameters π and μ. To identify π, we match the number of deaths predicted by the model with the observed deaths reported by the Colombian

Table 1

Calibrated Parameters

Parameter	Value
Calibrated with minimum distance algorithm	
π_1	1.5403×10^{-6}
π_2	1.0014×10^{-5}
π_3	0.9703
π_d	0.0061
π_r	0.7141
π_m	0.0185
μ	0.3783
Calibrated a la ERT	
A	19.2308
θ	0.0015
ϕ^1	0.8000
β	0.9992
κ	1.5000
ϵ	2.3897×10^{-5}

NOTE: ERT, Eichenbaum, Rebelo, and Trabandt (2020).
SOURCE: Authors' calculations.

health care authorities. We focus on the deaths to avoid problems associated with the underreporting of detected cases.

To identify μ, we match the weekly consumption predicted by the model to a proxy of weekly economic activity that we observe in almost real time. Specifically, we use electricity consumption, which historically roughly matches the economic cycle, with the understanding that electricity is both an input used in any type of consumption activity and difficult to substitute in the short run.

The model is calibrated over 31 weeks, starting during the second week of April 2020, which is eight weeks after the presumed beginning of the pandemic in Colombia. As occurs almost anywhere, there is very substantial underreporting of contagion, as most infections are asymptomatic. On the other hand, the Colombian health care system collects death data with accuracy. Therefore, instead of using reported cases, as in ERT, we calibrate the model matching the predicted and observed weekly deaths. Our proxy of consumption is the gap between observed weekly electricity consumption and a simulated trend estimated from historic data. The electricity consumption is reported daily by the national electricity system operator and aggregated into weekly data.

Our calibration algorithm minimizes the following loss function:

$$(11) \qquad \Omega = \sum_{t=1}^{t=31}\left[\omega\left(D_t - \hat{D}_t\right)^2 + (1-\omega)\left(E_t - \hat{E}_t\right)^2 \right],$$

where D_t and $\hat{D}_t$ are the observed and predicted weekly deaths, respectively. On the other hand, E_t and $\hat{E}_t$ are, respectively, the realized and predicted weekly electricity consumption gaps with respect to a scenario without the pandemic. The initial time period $t=1$ corresponds to the thirteenth week of 2020.

As mentioned previously, we use electricity consumption as a proxy to measure economic activity. However, to compute E_t we need a measure of the average electricity consumption that would have been observed in a scenario without a pandemic. We do so by projecting the trend of electricity consumption implied by the data observed up to March 2020.[3]

As for the weighting scalar, ω, we use a backtesting approach to choose the value that minimizes the average out-of-sample prediction error.[4]

3 RESULTS

3.1 Baseline

In Figures 1 and 2 we show the simulated and observed measures of the pandemic and the economic activity, respectively. The data correspond to the weekly number of COVID-19-related deaths and the weekly gap in electricity consumption. We calibrate two models. On the one hand, we calibrate our model with immune agents as specified above. On the other hand, we also calibrate a model with no immune individuals ($M_t = 0, \forall t$), which is equivalent to the standard macro-SIRD model a la ERT—that is, susceptible (S), infected (I), recovered (R), and dead (D). The contrast between both calibrations illustrates the contribution of our approach.

As shown in Figures 1 and 2, the macro-SIOD model is able to predict well the pattern of both variables. Relative to the observed deaths, the model predicts a later peak at a level of around 2,000 weekly deaths, which is slightly lower than the observed peak deaths. The number of observed deaths experienced a slight acceleration during September 2020 that the model cannot replicate.

Figure 1

Observed and Simulated Deaths

Deaths (thousands)
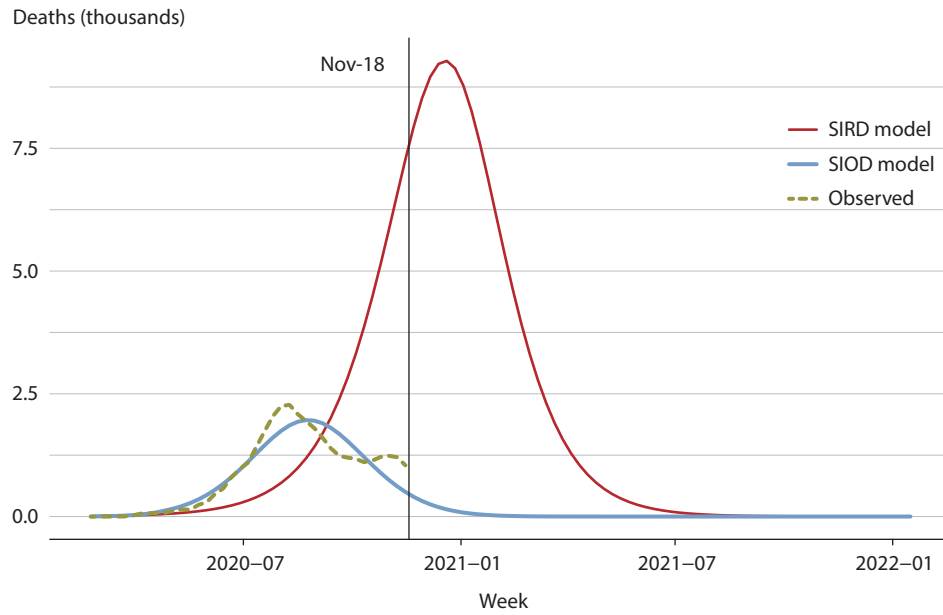
SOURCE: Authors' calculations based on data from the Instituto Nacional de Salud (INS).

Figure 2

Observed and Simulated Consumption Gap

Gap (percent)

SOURCE: Authors' calculations based on data from the Instituto Nacional de Salud (INS).

On the other hand, the model predicts well the collapse of economic activity observed during the strict quarantine that was in place during April 2020. The observed recovery afterward is much bumpier than the prediction of the model. In any case, the model predicts well the rapid recovery of economic activity, as measured by the electricity consumption gap.

The macro-SIOD model predicts that the COVID-19 pandemic will have been mostly over by the beginning of 2021. This prediction is relatively optimistic compared with standard epidemiological models that have been used to forecast the progress of the pandemic in Colombia. It should also be said that these standard models have consistently predicted much higher deaths than observed.

On the other hand, the macro-SIOD model predicts that economic activity will have converged to its long-run path by mid-2021. The model predicts that consumption will have fallen 4.5 percent below its long-run level in 2020 and have recovered almost fully in 2021. We should note that our calibration is based on electricity consumption, which is an imperfect measure of consumption and of economic activity in general. In particular, there are sectors of the economy that might be permanently affected, such as tourism and entertainment activities, that are not intensive in electricity use. It should not be a surprise, then, that the model predicts a full recovery, whereas a portion of the economy most probably will be underperforming for a long while.

In contrast to our macro-SIOD model, the calibrated macro-SIRD model has more difficulties matching the data. As shown in Figure 1, this model predicts a much later and higher peak in weekly deaths than what we observe in the data. The model predicts a total of almost 200,000 deaths, whereas the preferred macro-SIOD model predicts no more than 36,000 deaths. Standard epidemiological models, such as SIRD, commonly predict a much higher number of deaths than are observed.

As shown in Figure 2, the macro-SIRD model also has difficulties predicting the path of economic activity. The model replicates well the initial dip of consumption, but it then shows a second dip in early 2021. In this model, the second dip is a result of the relaxation of government restrictions, which increases substantially the risk of contagion and which in turn induces consumers to reduce consumption and labor.

The contrast between the calibrated SIRD and SIOD models suggests that the addition of immune/non-contagious agents allows the model to replicate the data much better than the standard model. The standard macro-SIRD model is unable to generate simultaneously reasonable predictions for both the number of deaths and economic activity, using our standard calibration methodology.

Because our model assumes that immune agents stay immune forever, it predicts that the pandemic ends after the first wave of contagion. It should be noted, though, that our model can be easily extended to allow immune agents to become susceptible again by calibrating an additional exogenous probability of immune agents reverting to susceptibility. Moreover, vaccinations with full or limited immunity can also be incorporated into the model. Since at this point it is not clear how immunity works, even with vaccines, we abstract from the problem and focus on a one-wave pandemic.

In the analysis that follows we use the preferred macro-SIOD model to evaluate counterfactual scenarios. We focus on the counterfactual behavior of the economy during 2020 and avoid using the model to predict the future, which will be affected by precisely how immunity evolves over time.

3.2 Counterfactual Analysis: The Impact of Government Restrictions and Individual Choices

We use the calibrated macro-SIOD model to evaluate the impact of imposed government measures and individual self-regulation on the pandemic and economic activity. Specifically, we simulate the model assuming that there are no government restrictions and that individuals do not account for the contagion risk when consuming or working. We call this last assumption "suboptimal consumption." We perform three counterfactual simulations under combinations of these counterfactual assumptions.

More specifically, each counterfactual simulation can be described as follows:

1. No government restrictions plus suboptimal consumption, $\mu_t = 0$, and consumers ignore the additional risk of contagion associated with c_t and n_t: This simulation assesses the joint effect of government restrictions and individual behavior. It provides an upper bound on deaths relative to the baseline model.

2. No government restrictions plus optimal consumption decisions, $\mu_t = 0$: In other words, there are no limits to consumption activities, and therefore individuals freely maximize their welfare, accounting for the risk of contagion.

3. Observed government restrictions plus suboptimal consumption: Individuals ignore the additional risk of contagion associated with c_t and n_t.

We describe each simulation below. A common feature of the simulations is that they all predict the pandemic will have been over by early 2021. As pointed out above, this is an optimistic forecast that is very robust in the model.

We should also reiterate that our consumption calculations are based on the use of electricity, which is not a precise measure of total consumption. In particular, activities intensive in personal interactions, such as dining at restaurants or meeting at entertainment venues, are less intensive in electricity use than are the production and consumption of, for example, manufactured goods. Therefore, the demand for electricity has shown a faster recovery than has the overall economy, and our model predicts a fast recovery as well, along with a full convergence to the long-run equilibrium path by the middle of 2021.

The Impact of Both Government Restrictions and Individual Behavior. We first simulate the model assuming that there are no restrictions, setting $\mu_t = 0$, $\forall t$, and assuming that consumers perceive the probability of contagion in (6) as not related to consumption and labor activities; that is, individuals believe $\pi_1 = \pi_2 = 0$ in (6).

In this fully unrestricted model, individuals behave as if the pandemic follows the standard epidemiological model that assumes an exogenous probability of contagion. Nevertheless, the actual probability of contagion in (1) is still affected by the behavior of individuals. Under our assumption, individuals are irrational in the sense that they believe the contagion risk is exogenous and given solely by π_3. Therefore, in this model the effect of the pandemic on economic activity is very low and driven only by the number of individuals who die, which is a small share of total population.

The results of the simulation are shown in Figures 3 and 4. The predicted number of deaths in this simulation is 50,506, which is almost 42 percent higher than in the baseline. Recall that this number of deaths would have been the result if the government had imposed no restrictions and if

Figure 3

Simulated Deaths: Baseline vs. No Restrictions and Suboptimal Consumption

Deaths (thousands)

SOURCE: Model's predictions.

Figure 4

Simulated Consumption Gap: Baseline vs. No Restrictions and Suboptimal Consumption

Gap (percent)

SOURCE: Authors' calculations based on data from the Colombian Electricity Independent System Operator (XM).

individuals had not changed their behavior endogenously. In that sense, this figure is the upper bound in the number of deaths according to the calibrated model.

Notice that the consumption path is almost constant, which is a reflection of the fact that behavior is not affected by the pandemic in this simulation.[5] The baseline consumption in 2020 is 4.7 percent lower than this unrestricted consumption level. This figure is a rough estimate of the economic cost of the pandemic in the model.

The Role of Government Restrictions. We now simulate the model assuming that there are no restrictions but that individuals fully account for the effects of their behavior on the risk of contagion. In other words, we set $\mu_t = 0$, $\forall t$ and keep the remaining parameters of the model as in the baseline simulation. The results of this simulation are shown in Figures 5 and 6.

As shown in Figure 5 and as expected, the imposed quarantine did have a substantial effect on the number of deaths. Without the restrictions, the model predicts a total of 45,654 deaths by the end of the pandemic. The model implies that the excess deaths would have occurred mostly around the peak. The restrictions delayed the peak for several weeks and decreased its level by around 1,000 deaths per week.

As shown in Figure 6, the model without government restrictions shows a much smaller dip in consumption than observed, which coincides in time with the predicted peak in deaths. In the model, this decrease in consumption is a result of consumers' efforts to avoid contagion. In other words, the restrictions had an immediate and substantial effect on economic activity.

The Role of Individual Behavior. Finally, we isolate the impact of individual efforts to self-regulate their behavior on both contagion and economic activity. We fix government restrictions as in the baseline calibration but set the perceived probability of contagion in (6) equal to zero; that is, $\pi_1 = \pi_2 = 0$. As explained above, in this model, individuals believe that contagion risk is given by π_3. Because contagion is perceived to be unaffected by behavior, the simulated dip in consumption is entirely a result of government restrictions.

We show the results of this simulation in Figures 7 and 8: There are around 2,800 more deaths in this scenario than in the baseline simulation. Compared with the results of the previous simulation, the model suggests that individual self-regulation had less of an effect on deaths than did government restrictions. The effect on economic activity is an increase of around 1.4 percent relative to the 2020 baseline value. In other words, the change in individual behavior explains a relatively small portion of the decrease in economic activity.

We show a summary of the results of these simulations in Table 2, including consumption and deaths in 2020 and 2021 for the baseline and each simulation. As pointed out above, without government restrictions and agents' behavior, the total number of deaths would have been 50,506, which is 42 percent higher than in the baseline simulation. Consumption in 2020 would have been 4.7 percent higher than in the baseline simulation and would have grown less than 1 percent in 2021. These figures are a measure of the total cost of the pandemic, in terms of both deaths and economic activity.

Without government restrictions, agents' behavior would have resulted in 45,458 deaths. Therefore, government restrictions saved 10,060 lives, which is 28 percent of baseline deaths. In this scenario, consumption would have been 3 percent higher in 2020 and then would have fully recovered in 2021.

Figure 5

Simulated Deaths: Baseline vs. No Restrictions and Fully Optimal Consumption

Deaths (thousands)

SOURCE: Model's predictions.

Figure 6

Simulated Consumption Gap: Baseline vs. No Restrictions and Fully Optimal Consumption

Gap (percent)

SOURCE: Authors' calculations based on data from the Colombian Electricity Independent System Operator (XM).

Figure 7

Simulated Deaths: Baseline vs. Suboptimal Consumption with Restrictions

Deaths (thousands)

SOURCE: Model's predictions.

Figure 8

Simulated Consumption Gap: Baseline vs. Suboptimal Consumption with Restrictions

Gap (percent)

SOURCE: Authors' calculations based on data from the Colombian Electricity Independent System Operator (XM).

The role of individual behavior is more limited. Without the rational response of individuals to the risk of contagion, the number of deaths would have been 38,458, which is 2,864 more than in the baseline simulation. Therefore, individual behavior saved 8 percent of deaths, relative to the baseline. In this scenario, 2020 consumption would have been only 1.4 percent higher than in the baseline simulation and then would have almost fully recovered in 2021.

These results imply that the combination of policy and individual behavior saved around 42 percent of baseline deaths, with an economic cost of around 4.7 percent of consumption in 2020. The simulations imply that government restrictions had a bigger impact than the endogenous changes in individual behavior.

4 CONCLUSION

We have calibrated a model of economic behavior during the COVID-19 pandemic, as it applies to the Colombian economy. Our model incorporates an "immune" type of agent that better explains the data than do standard epidemiological models. In our model, the pandemic falls rapidly and disappears during early 2021. Consumption falls substantially during 2020 but recovers fully by mid-2021. It also implies that government restrictions and consumers' self-regulation helped to avert around 15,000 deaths, or around 42 percent of baseline deaths. Government restrictions account for more than 67 percent of this effect.

The model focuses on the first wave of the pandemic and is able to reproduce only this first wave. However, a shortcoming of the model is the assumption that immune agents stay immune forever. To generate additional waves of contagion, the model can be extended to allow for limited immunity so that immune agents become susceptible again after a period of time, or with some probability. Immunity can also be modeled to incorporate vaccinations. Given the uncertainty that surrounds the evolution of immunity to this virus over time, we leave these issues for future research. ∎

Table 2

Summary Results of Simulated Scenarios

Baseline	
Deaths up to Dec 2020	35,311
Deaths up to Jun 2021	35,594
Deaths total	35,594
Consumption 2020 (trillions COP)	979.6
Consumption 2021 (trillions COP)	1,027.3
No restrictions plus suboptimal consumption	
Deaths up to Dec 2020	50,323
Deaths up to Jun 2021	50,506
Deaths total	50,506
Consumption 2020 (trillions COP)	1,027.9
Consumption 2021 (trillions COP)	1,027.9
No government restrictions	
Deaths up to Dec 2020	45,458
Deaths up to Jun 2021	45,654
Deaths total	45,654
Consumption 2020 (trillions COP)	1,009.6
Consumption 2021 (trillions COP)	1,027.9
Suboptimal consumption	
Deaths up to Dec 2020	38,172
Deaths up to Jun 2021	38,458
Deaths total	38,458
Consumption 2020 (trillions COP)	993.2
Consumption 2021 (trillions COP)	1,027.3

NOTE: COP, Colombian peso.
SOURCE: Authors' calculations.

APPENDIX

In this appendix, we show the full set of equilibrium conditions of the model. The definition of the variables and their transitions over time are defined in the body of the article. The additional equilibrium conditions of the model imply that all agents maximize their expected lifetime utility, firms maximize profits, the government balances its budget, and both labor and goods markets clear.

First-order conditions for c and n: For $j = S$ (susceptible agents),

$$u_1\left(c_t^S, n_t^S\right) - \left(1 + \mu_t\right)\lambda_{bt}^S + \lambda_{\tau t}\left(1 - \pi_m\right)\pi_1\left(I_t C_t^I\right) = 0 \text{ and}$$
$$u_2\left(c_t^S, n_t^S\right) + w_t\lambda_{bt}^S + \lambda_{\tau t}\left(1 - \pi_m\right)\pi_1\left(I_t N_t^I\right) = 0.$$

First-order conditions for c and n: For $j = I$ (infected agents),

$$u_1\left(c_t^I, n_t^I\right) = \lambda_{bt}^I\left(1 + \mu_t\right) \text{ and}$$
$$u_2\left(c_t^I, n_t^I\right) = -\phi^I w_t\lambda_{bt}^I.$$

First-order conditions for c and n: For $j = O$ (surviving agents),

$$u_1\left(c_t^O, n_t^O\right) = \lambda_{bt}^O\left(1 + \mu_t\right) \text{ and}$$
$$u_2\left(c_t^O, n_t^O\right) = -w_t\lambda_{bt}^O.$$

The first-order condition for τ_t is

$$\beta\left(1 - \pi_m\right)\left(U_{t+1}^I - U_{t+1}^S\right) - \lambda_{\tau t} = 0.$$

The government budget constraint implies that the government balances its budget; that is,

$$\mu_t\left(S_t c_t^S + I_t c_t^I + O_t c_t^O\right) = \Gamma_t\left(S_t + I_t + O_t\right).$$

Finally, market-clearing conditions guarantee that both labor and goods markets clear; that is,

$$\begin{aligned} S_t C_t^S + I_t C_t^I + O_t C_t^O &= A N_t \\ S_t N_t^S + I_t N_t^I \phi^I + O_t N_t^O &= N_t \end{aligned}.$$

NOTES

[1] As pointed out by a referee, immunity arises independently of agents' decisions. If immunity were a consequence of interactions, it would be equivalent to infection and the model would revert to the standard SIRD model: susceptible (*S*), infected (*I*), recovered (*R*), and dead (*D*).

[2] An alternative specification would simply assume an exogenous number of immune agents. Our current specification recognizes the possibility that agents can move in and out of immunity and would easily accommodate vaccinations with limited effectiveness.

[3] More specifically, we estimate such a trend fitting an autoregressive integrated moving average (ARIMA) model with both monthly fixed effects and a time trend and then project the implied, predicted mean starting from March 2020.

[4] Our backtesting approach uses a moving window of 10 weeks starting from the third week of March 2020. After standardizing scales—that is, dividing each series by its own sample standard deviation—we find $\omega = 0.55$.

[5] It is worth noting that the slight drop in consumption of this counterfactual is due to the lower productivity faced by infected workers, which directly affects consumption in the equilibrium.

REFERENCES

Acemoglu, D.; Chernozhukov, V.; Werning, I. and Whinston, M.D. "Optimal Targeted Lockdowns in a Multi-Group SIR Model." NBER Working Paper 27102, National Bureau of Economic Research, 2020; https://doi.org/10.3386/w27102.

Adam, D. "A Guide to *R*—the Pandemic's Misunderstood Metric." *Nature*, 2020, *583*(7816) pp. 346-48; https://doi.org/10.1038/d41586-020-02009-w.

Alvarez, F.E.; Argente, D. and Lippi, F. "A Simple Planning Problem for COVID-19 Lockdown." NBER Working Paper 26981, National Bureau of Economic Research, 2020; https://doi.org/10.3386/w26981.

Atkeson, A. "What Will be the Economic Impact of COVID-19 in the U.S.? Rough Estimates of Disease Scenarios." Staff Report 595, Federal Reserve Bank of Minneapolis, 2020; https://doi.org/10.21034/sr.595.

Berger, D.W.; Herkenhoff, K.F. and Mongey, S. "An SEIR Infectious Disease Model with Testing and Conditional Quarantine." Staff Report 597, Federal Reserve Bank of Minneapolis, 2020; https://doi.org/10.21034/sr.597.

Doshi, P. "COVID-19: Do Many People Have Pre-Existing Immunity?" *BMJ*, 2020, *370*; https://doi.org/10.1136/bmj.m3563.

Eichenbaum, M.S.; Rebelo, S. and Trabandt, M. "The Macroeconomics of Epidemics." NBER Working Paper 26882, National Bureau of Economic Research, 2020; https://doi.org/10.3386/w26882.

Kermack, W.O. and McKendrick, A.G. "A Contribution to the Mathematical Theory of Epidemics." *Royal Society*, 1927, *115*(772) pp. 700-21; https://doi.org/10.1098/rspa.1927.0118.