

Expectations, Infections, and Economic Activity

Martin Eichenbaum

Northwestern University and National Bureau of Economic Research

Miguel Godinho de Matos

Universidade Católica Portuguesa

Francisco Lima

Universidade de Lisboa

Sergio Rebelo

*Northwestern University, National Bureau of Economic Research,
and Center for Economic Policy Research*

Mathias Trabandt

Goethe University Frankfurt

This paper develops a quantitative theory of how people weigh the risks
of infections against the benefits of engaging in social interactions that

A previous version was titled “How Do People Respond to Small Probability Events with Large, Negative Consequences?” We thank Laura Murphy and Joao Guerreiro for excellent research assistance and the editor, Andrew Atkeson, for his extensive comments. In addition, we thank Eran Yashiv and two anonymous referees for their feedback. Miguel

Electronically published July 18, 2024

Journal of Political Economy, volume 132, number 8, August 2024.

© 2024 The University of Chicago. This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), which permits non-commercial reuse of the work with attribution. For commercial use, contact journalpermissions@press.uchicago.edu. Published by The University of Chicago Press.
<https://doi.org/10.1086/729449>

contribute to the spread of infectious diseases. Our framework takes into account the effects of public policies and private behavior on the spread of the disease. We evaluate the model using a novel micro panel dataset on consumption expenditures of young and older people across the first three waves of COVID-19 in Portugal. Our model highlights the critical role of expectations in shaping how human behavior influences the dynamics of epidemics.

I. Introduction

It is now widely recognized that human behavior influences the dynamics of epidemics. But how should we incorporate the impact of human actions into quantitative epidemiological models?

In our view, a successful approach requires a quantitative theory of how people weigh the risks of infections against the benefits of engaging in social interactions that contribute to the spread of infectious diseases. The resulting model should also account for the interrelated yet distinct effects of public policies and private behavior on the spread of the disease.

We develop such a model and evaluate its plausibility using a novel micro dataset on consumption expenditures in Portugal. In so doing, we encounter two key challenges. The first is to account for the cross-sectional consumption response of individuals of different ages at a given point in time. This response is consistent with a full-information rational expectations (FIRE) model, in which the old rationally fear that they are more likely than the young to die from an infectious disease. The second challenge is to account for the time series response of consumption of the young and old across the first three waves of COVID. The consumption responses in the first and third waves are similar, but deaths per capita were much larger in the third wave than in the first. At the same time, government-imposed containment measures were similar in the first and third waves. These observations are inconsistent with a simple FIRE model.

We develop a quantitative model that meets the challenges discussed above. In so doing, we face a set of difficulties likely to be encountered by any researcher modeling an epidemic in a specific location and the observed behavioral responses to disease threats over extended time horizons. These difficulties include how to (i) model the evolution over time of individual beliefs about the risks presented by a new disease, (ii) model

Godinho de Matos was funded by the Portuguese Foundation for Science and Technology grant UID/GES/00407/2020. The views expressed in this article are those of the authors and do not necessarily reflect the official views of authors' institutions. The data we use are part of Statistics Portugal's project of making administrative data from the Portuguese tax authority and other public agencies available for statistical production and research. For information on how to access the data, see the Statistics Portugal website.

the evolution over time of beliefs that individuals have about the persistence of this risk, (iii) model individual beliefs about the impact of the actions that they can take to mitigate the risks of infection, (iv) isolate the separate roles of risk aversion (uncertainty over health outcomes) and intertemporal substitution in shaping behavior, (v) model the extent to which individuals care about dying (the value of bequests versus the value of living), (vi) reconcile the large short-run impact of COVID on consumption expenditures with the small corresponding impact of the secular decline in mortality from infectious diseases, and (vii) separate the quantitative impacts of public policy, such as lockdowns and private behavior, on economic and disease outcomes. We model lockdowns as a wedge in the consumers' utility and assume that the magnitude of this wedge is proportional to an index of the severity of lockdowns in Portugal constructed by an external source. We estimate and judge the importance of lockdowns by the impact of changes in the proportionality factor on model fit. Our approach to modeling containment is equivalent, in many contexts, to an absolute prohibition on purchases of certain goods (see app. B.3; apps. A–C are available online).

Our answers to the previous questions are not definitive. But we hope our analysis is a useful step in quantifying the different forces at work that any behavioral epidemiological model will have to incorporate.

Our paper makes five specific contributions to the literature. First, we use micro data to document the empirical response of consumption expenditures to COVID by people of different ages, comorbidity statuses, and incomes.

Second, we estimate our structural model using micro data on consumption expenditures. The key parameters that we estimate include old and young people's prior beliefs about case fatality rates and the speed with which they change their views. We find that all people had pessimistic prior beliefs about case fatality rates but learned the actual case fatality rates over time.

Third, we highlight the importance of deviations from full-information rational expectations in accounting for the empirical response of consumption to the COVID epidemic.

Fourth, we use the empirically validated model to assess how much people of different ages and incomes would be willing to pay to avoid the epidemic. Naturally, people's expectations about mortality rates play a crucial role in their willingness to pay. We also explore the distinct roles of intertemporal substitution and risk aversion in determining the willingness to pay.

Fifth, we suggest a way to reconcile the large short-run and small long-run effects on consumption of changes in mortality rates associated with contagious diseases. Our suggestion highlights the critical role of expectations about case fatality rates in such a reconciliation. To keep our analysis

tractable, we abstract from long-run supply issues possibly arising from changes in fertility and education decisions.

Our paper is organized as follows. Section II briefly reviews the related literature. Section III describes our data. Section IV contains our empirical results. Section V presents a simple model used to develop intuition about the mechanisms at work in our quantitative model. Section VI describes the quantitative model and estimation procedure. Section VII summarizes our estimation results. Section VIII contains a general equilibrium model of endemic COVID. This model extends our partial equilibrium analysis along three dimensions. First, we embed it in a general equilibrium framework with endogenous labor supply and capital accumulation. Second, we allow for vaccination. Third, we modify the epidemiology assumptions so that people who have natural immunity or are vaccinated lose their immunity over time. We conclude in section IX.

II. Related Literature

There is, by now, an extensive literature on the macroeconomic impact of epidemics. Examples include Krueger, Uhlig, and Xie (2020), Toxvaerd (2020), Alvarez, Argente, and Lippi (2021), Eichenbaum, Rebelo, and Trabandt (2021, 2022a, 2022b), Faria-e-Castro (2021), Farboodi, Jarosch, and Shimer (2021), Jones, Philippon, and Venkateswaran (2021), Guerrieri et al. (2022), and Piguillem and Shi (2022). There is also a sizable epidemiology literature on the interaction between COVID and age. Examples include Dessie and Zewotir (2021), Doerre and Doblhammer (2022), and Sorensen et al. (2022). We do not attempt to survey these literatures here. Instead, we discuss the papers most closely related to ours in the sense that they study the impact of age on people's consumption behavior. The three key papers are as follows.

Glover et al. (2020) analyze a two-sector model (essential and luxury) with young workers and retirees. The epidemic creates significant distributional effects because the luxury sector contracts more than the essential sector. In addition, containment measures redistribute welfare from the young to the old. The old benefit from the reduced risk of infection produced by containment, while the young suffer the adverse employment consequences.

Brotherhood et al. (2021) use a calibrated model of the pandemic that features age heterogeneity and individual choice, allowing agents to choose rationally how much social distancing to undertake, considering future infection risk and prospects for vaccine arrival.

Acemoglu et al. (2021) study targeted lockdowns in a multigroup susceptible-infected-removed (SIR) model where infection, hospitalization, and fatality rates vary between groups, in particular, between the young, middle-aged, and old.

III. Data

Our dataset comes from Statistics Portugal, the national statistical authority. A software system called e-fatura—which the Portuguese government adopted in 2013 to reduce tax evasion—generates the data. The decree law 198/2012 published on August 24, 2012, requires firms to report their invoice data electronically. This decree covers all individuals or legal entities with headquarters, stable establishment, or tax domicile in Portuguese territory that conduct operations subject to value-added tax (VAT). Durable goods purchases—such as cars, refrigerators, and televisions—are included in our dataset because they are subject to VAT. However, we cannot separate purchases of durable and nondurable goods because we cannot access itemized invoices that specify the nature of the goods purchased.

Goods and services exempt from VAT are excluded from the data.¹ The most important exempt categories are health services provided by medical doctors, childcare services provided by kindergartens, residential homes, day centers for the elderly, rent and property investments, and services provided by nonprofit organizations that operate facilities for art, sports, or recreation activities. Our data cover approximately 75% of the per capita consumption expenses included in the national income accounts.

Our data include anonymized information for 500,000 Portuguese people randomly sampled from a set of 6.3 million people who meet two criteria. First, they were at least 20 years old in 2020. Second, they filed income taxes as Portuguese residents in 2017. The dataset includes a person's age, income bracket, and gender. In addition, for a subset of people, the data include education and occupation in 2017.

For every person in our sample, we construct total monthly consumption expenditures using the electronic receipts that firms provide to the tax authority as part of their VAT reporting. Each receipt is matched to a particular person using their anonymized fiscal number. We also compute individual pharmacy expenditures, which we use as a proxy for comorbidity.

Portuguese consumers have four incentives to include their fiscal number in expenditure receipts. First, they can deduct from their income taxes, up to a limit, expenditures on health, education, lodging, nursing homes, and general household spending. Second, the government rebates 15% of the VAT from documented expenditures on public transportation passes, lodging, restaurants, and automobile and motorcycle shops. Third, for every €10 of reported spending, consumers receive a coupon for a weekly lottery in which the prize is a 1-year treasury bond with a face value of €35,000. Fourth, the law obliges consumers to

¹ See article 9 of the VAT code for an exhaustive list.

request invoices for all purchases of goods and services. Consumers who fail to comply are subject to fines ranging from €75 to €2,300.

Young people might have purchased goods and services for their parents and reported them under their own fiscal number. But it is in general not in their interest to do so because there are caps on the VAT rebates that taxpayers can receive and on the VAT expenses that taxpayers can deduct from their income taxes.

The data include online purchases from Portuguese businesses but excludes online purchases from foreign companies. The latter types of purchases are likely to be small and not negatively affected by COVID. Since young people are more likely to engage in such purchases, including them would likely strengthen the result (documented below) that older people cut their consumption by more than young people.

We exclude from the sample in a given month people who do not have any receipts associated with their fiscal number for that month. We also remove from the sample 21,814 people who were unemployed or inactive in 2017. These people are unlikely to pay taxes, so they have less incentive to include their fiscal number in receipts. Finally, we dropped all persons older than age 80 from the sample because their expenditure patterns suggest that many of them live in nursing homes. We also exclude people younger than age 20 because they make few independent consumption decisions. The resulting dataset contains 421,337 people and 12,218,773 person-month observations aggregated over 97,363,250 buyer-seller pairs.

We identify two groups in our sample whose incomes are likely to have been relatively unaffected by the COVID recession: public servants (58,598 people) and retirees (93,839 people). These groups overlap because we do not exclude retirees from the population of public servants. There are roughly 22,000 retired public servants in our sample.²

Our sample covers the period from January 2018 to April 2021. We end our sample in April 2021 for two reasons. First, vaccines became available to the general population after April 2021. Before April, only the elderly and people with comorbidities were vaccinated first. Second, according to data from the Global Initiative on Sharing All Influenza Data (GISAID), there were no reported cases of the delta variant, which was arguably more contagious than previous variants.³

Table 6 (tables 6–16 are available online) reports descriptive statistics for monthly expenses net of VAT. For public servants, the average per capita monthly expenditure on consumption goods and services is €687.8, of

² In 2011, Portugal entered into an adjustment program with the International Monetary Fund, the European Central Bank, and the European Commission (for a discussion, see Eichenbaum, Rebelo, and Resende 2017). This reduction led to a large increase in the number of retired public servants.

³ Data are from <https://covariants.org>.

which €25.6 is spent on pharmacy items. These expenditures are roughly similar for the sample of the population as a whole: the average per capita monthly expenditure on consumption goods and services is €629.3, of which €17.9 is spent on pharmacy items. Retirees have lower levels of overall expenditure. They spend, on average, €437.8 on consumption goods and services, of which €24.3 is spent on pharmacy items.

Table 7 reports the same statistics as table 6 broken down by income and age groups. Income groups are based on the 2017 income tax brackets used by Portugal’s Internal Revenue Service. We group people according to their ages so that they have similar COVID case fatality rates. Our estimates of this risk are based on the statistics reported by the Portuguese Health Authority (Direção-Geral da Saúde) on July 28, 2020. Table 1 displays case fatality rates (the ratio of COVID deaths to people infected) by age cohort for Portugal. Two key results emerge from table 1. First, people aged 20–49 all have low case fatality rates. Second, case fatality rates rise nonlinearly with age for people older than age 50.

IV. Empirical Results

This section has two parts. In section IV.A, we provide an overview of the evolution of the epidemic in Portugal and the government’s containment measures. We also discuss the evolution of per capita consumption expenditures in our sample. In section IV.B, we present formal econometric evidence of how COVID impacted the consumption expenditures of people of different ages and comorbidity conditions.

A. The Epidemic in Portugal

Figure 1 depicts the weekly time series of infected people and COVID deaths in Portugal. We refer to March 2020 through April 2021 as the epidemic dates. There were three waves of COVID deaths during this period. The peaks of these waves occur in April 2020, December 2020,

TABLE 1
COVID INFECTION FATALITY RATES (Averages, May 14–June 14, 2020)

Age Group	Infected	Deceased	Infection Fatality Rate (%)
0–9	672	0	.0
0–19	1,085	0	.0
20–29	4,245	1.5	.03
30–39	4,869	.6	.01
40–49	5,420	15.3	.28
50–59	5,336	43.6	.82
60–69	3,519	122.1	3.5
70–79	2,576	265.9	10.3
≥80	4,522	926	20.5

NOTE.—Rates are computed with data from the Portuguese Health Authority.

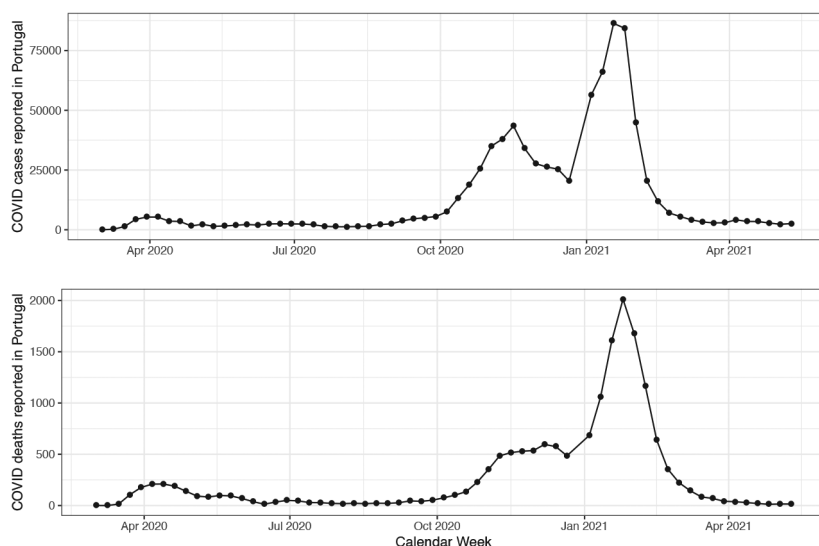


FIG. 1.—COVID-19 cases and deaths reported by Portuguese Health Authority (May 20, 2021).

and January 2021. The broad pattern of COVID cases is consistent with the facts documented by Atkeson, Kopecky, and Zha (2020) for a cross section of countries.

The vaccination campaign started on January 8, 2021. The initial campaign focused on people over age 80 with comorbidities. Vaccination of the general population began on April 23, 2021, very close to the end of our sample (April 30, 2021).

Over the period from March 2020 to April 2021, the government implemented various containment measures. These measures vary in intensity and sectoral coverage. For concreteness, we summarize the severity of these measures using an index of the full or partial closing of nonessential shops, restaurants, and cafés.⁴ Figure 2 displays this containment index. Containment rose quickly in mid-March 2020 and started to decline at the beginning of May 2020. It then dropped to low levels in the summer of 2020. In mid-November 2020, containment was partially reimposed in response to the second wave. The third epidemic wave led to the strengthening of containment measures from January to March 2021. As the number of infections waned, containment measures were eased. Note that the peak containment rates are the same in the first and third waves.

⁴ To construct this index, we use data from <https://ourworldindata.org> and <https://dre.pt/legislacao-covid-19-upo>. We attribute the values 1, 0.5, 2/7, and 0 to full closing, partial closing, closing on weekends, and open. The containment index is the average of the indexes for nonessential shops and restaurants and cafés.

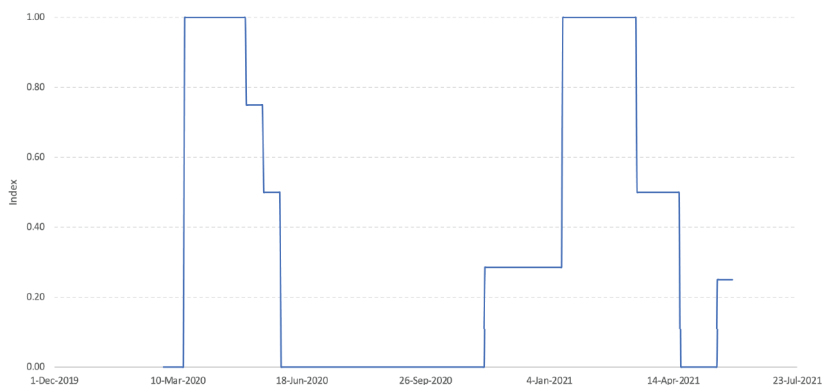


FIG. 2.—Severity of COVID-19 containment measures over time.

Figure A.1 (figs. A.1–B.10 are available online) depicts the average logarithm of public servants' monthly consumption expenditures from January 2018 to April 2021. Three features emerge from this figure. First, there are pronounced drops in consumption around the peak months of the first and third waves. There is a more muted decline in consumption during the months around the peak of the second wave. Second, there is a clear seasonal pattern in the pre-COVID sample. This pattern is similar in 2018 and 2019. Third, per capita spending was growing before the COVID shock. Our econometric procedure considers the latter two features in creating a counterfactual for what spending would have been in 2020 absent the COVID shock. We estimate a seasonal effect and time trend for each age and income group using data from January 2018 to February 2020.

B. Age and the Impact of COVID on Consumer Expenditures

Our empirical specification focuses on the differential consumption response by people of different ages. This specification is given by

$$\ln(Expenses_{it}) = \Lambda \times Year_t + \sum_{s=Feb}^{Dec} \lambda_s \mathbf{1}\{Month_t = s\} + \theta_i + \Psi_{it} + \epsilon_{it}$$

$$+ \sum_{m=Mar\ 2020}^{Apr\ 2021} \Delta_m After_t \times \mathbf{1}\{Date_t = m\} +$$

$$\sum_{m=Mar\ 2020}^{Apr\ 2021} \sum_{g \in AgeGroup \setminus [20;49]} \delta_{mg} After_t \times \mathbf{1}\{Date_t = m\} \times \mathbf{1}\{AgeGroup_i = g\}.$$

Subscripts i and t denote person i and calendar month t , respectively. The coefficient Λ represents a linear growth trend in consumption expenditures.

$Year_t$ is a variable that takes the value $1 + t$ for year $2018 + t$ for $t = 0, 1, 2, 3$. The coefficients λ_s control for seasonality in consumption. The vector Ψ_{it} includes interaction terms that allow seasonal effects to vary with individual characteristics (age, income bracket, gender, education, and occupation). The coefficients θ_i denote time-invariant individual fixed effects. $After_t$ is a dummy variable equal to 1 during the epidemic dates (beginning March 2020). The coefficients Δ_m represent the change in spending for people in the reference group (aged 20–49) during the epidemic date d . The coefficient δ_{mg} measures the additional change in spending for age group g in epidemic date m .⁵ The variable ϵ_{it} is an idiosyncratic error term. As long as the inflation rate for the consumption baskets of different age cohorts is the same, any inflation effects cancel out from the difference in nominal responses, and we are left with the real differential response. We estimate equation (1) using a fixed effects estimator and cluster standard errors by person, as suggested in Bertrand, Duflo, and Mullainathan (2004).⁶

Column 4 of table 14 reports our parameter estimates. Figure 3 displays our estimates of the impact of COVID on consumption expenditures of different age groups (Δ_m for the reference group and $\Delta_m + \delta_{mg}$ for the other groups) obtained from estimating equation (1). The bars around the point estimates represent 95% confidence intervals. Our key findings are as follows. First, all consumers reduced their expenditures during the three waves of the epidemic. Second, older people cut their expenditures by much more than younger people. The nonlinear effect of age on consumer expenditures mirrors the nonlinear dependency of case fatality rates on age. Third, the decline in consumption for each age group was similar in the first and third waves.

C. *The Response of People with Different Income*

The economic model discussed in section VI implies that high-income people cut their expenditures by more than low-income people to reduce the risk of infection. According to the model's logic, rich people have more to lose from becoming infected than poor people. Since older people might have a higher income than younger people, the results reported in section IV.B might conflate the effect of age and income.

Table 15 reports our parameter estimates. Figure 4 displays our estimates of the impact of COVID on consumption expenditures of different age groups (Δ_m for the reference group and $\Delta_m + \delta_{mg}$ for the other groups) obtained from estimating equation (1) for separate income groups. Two key results emerge from this figure. First, our results about

⁵ We keep age groups constant on the basis of a person's age in 2020.

⁶ Because of our large sample size, we estimate the fixed effects models using the method of alternating projections implemented in R by Gaure (2013) and in Stata by Guimarães and Portugal (2010) and Correia (2016).

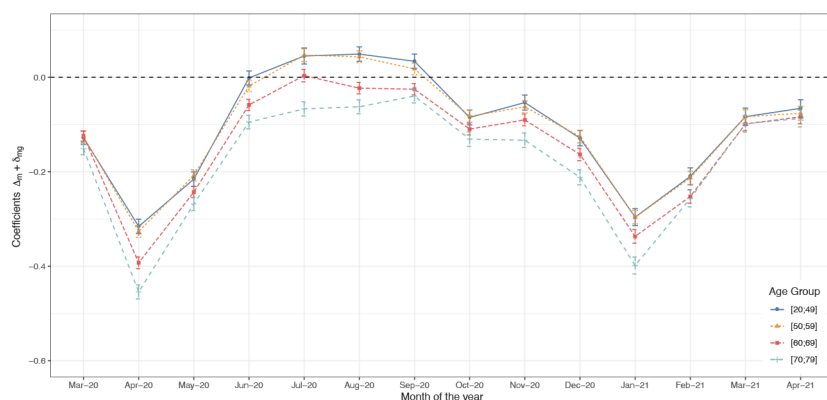


FIG. 3.—Changes in expenditures of public servants during epidemic relative to counterfactual without COVID.

the impact of age on consumption expenditures are very robust to controlling for income. Older people cut their expenditures by much more than younger people for all income groups. Second, when we control for age, high-income people reduce their consumption by more than low-income people.

The finding that expenditure cuts are an increasing function of income complements the evidence in Carvalho et al. (2020) and Chetty et al. (2020), which relies on home address zip codes to proxy for income.

D. Robustness

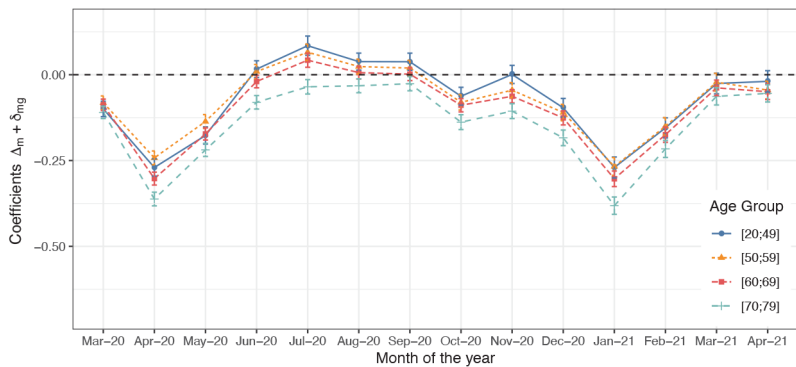
In appendix A, we report the results of six robustness checks. First, we provide evidence in favor of the assumption that the seasonal effects for January 2020 through April 2021 are the same as for the 2018–19 period.

Second, we redo our benchmark analysis, allowing for different monthly expenditure time trends for each age cohort. We find a similar pattern for the impact of age on the response of expenditures to the COVID shock.

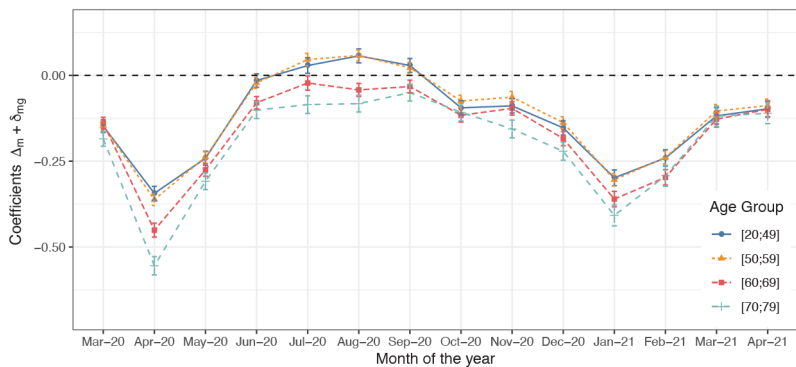
Third, we redo our empirical analysis for retirees instead of public servants. Retirees are another group whose income is likely to have remained relatively stable during the epidemic. Our results are similar to those that we obtain for public servants. We find that conditioning on age, the consumption expenditures of civil servants and retirees respond similarly to COVID.

Fourth, we find that our results are robust to running regression (1) using the year-on-year growth rate ($\ln(\text{Expenses}_{it}/\text{Expenses}_{it-12})$) instead of the log level of expenditures as the dependent variable.

A Income $\leq 20,261$



B $20,261 < \text{Income} \leq 40,522$



C Income $> 40,522$

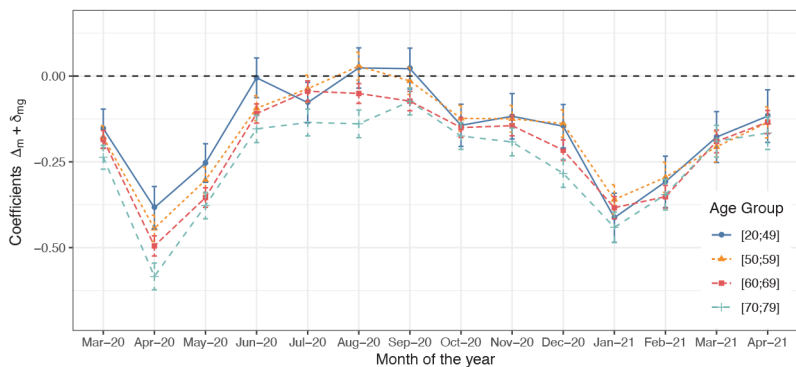


FIG. 4.—Changes in expenditures of public servants in different income groups during epidemic relative to counterfactual without COVID.

Fifth, we study a potential reason why the consumption expenditures of old and young people responded differently to COVID: these groups purchase different goods and services that were differentially affected by lockdowns. To investigate this possibility, we estimate the change in consumption expenditures for different age groups in sectors of the economy that were least affected by lockdowns. We base this sector classification on the information reported in the appendix to law 78-A/2020, approved September 29, 2020. Figure A.4, which is the analog to figure 3, presents our results. Two features are worth noting. First, all groups cut their consumption expenditures by about the same amount in the epidemic's first and third waves. Second, the old cut their consumption by more than the young in the epidemic's first, second, and third waves.

Sixth, we redo our analysis excluding two sectors where adaptations were most likely to have reduced the risk of infections: restaurants (people could order takeout instead of eating at the restaurant) and supermarkets (people could ask for delivery instead of going to the store). Figures A.5 and A.6 show that our results are robust to excluding these two expenditure categories.

Finally, we use data on expenditures on pharmaceutical drugs to investigate the effect of comorbidities that increase the risk of dying from COVID. We find that people with comorbidities cut their consumption more than those without comorbidities.

V. A Simple Model of Mortality Risk and Consumption Decisions

In this section, we consider a simple two-period model to develop intuition about the key features of our quantitative model presented in section VI. Consistent with the latter, we make three assumptions. First, the probability of dying depends on current consumption. Second, people derive utility from leaving a bequest when they die. Third, people's utility has the recursive form proposed by Kreps and Porteus (1978), Weil (1989), and Epstein and Zin (1991). These preferences allow us to study the different roles that risk aversion and intertemporal substitution play in our model.

In the first period of their life, a person receives an endowment, y , which they can consume in period 1 (c_1) or period 2 (c_2). Their resource constraint is

$$y = c_1 + c_2. \quad (2)$$

The period 1 utility is given by the following version of equation (8) in section VI:

$$U_1(y) = \left\{ (1 - \beta)c_1^{1-\rho} + \beta[E(U_2^{1-\alpha})]^{(1-\rho)/(1-\alpha)} \right\}^{1/(1-\rho)}. \quad (3)$$

The certainty equivalent of period 2 utility is

$$[E(U_2^{1-\alpha})]^{1/(1-\alpha)} = \{[1 - \delta(c_1)]c_2^{1-\alpha} + \delta(c_1)(\omega_0 + \omega_1 b^\mu)^{1-\alpha}\}^{1/(1-\alpha)},$$

where $\delta(c_1)$ is the probability of dying before consuming in period 2. To capture the basic mechanism at work in our epidemiological model, we assume that $\delta(c_1)$ is an increasing linear function of c_1 :

$$\delta(c_1) = \Gamma_0 + \Gamma_1 c_1, \quad (4)$$

where Γ_0 and Γ_1 are positive constants. A person who survives in period 2 consumes c_2 . A person who dies leaves their planned consumption, c_2 , as a bequest: $b = c_2$. The representative person chooses c_1 , c_2 , and b to maximize (3) subject to (2), (4), and $b = c_2$.

To derive the first-order conditions, it is useful to consider the following monotonic transformation of the Epstein-Zin utility function: $V_1 = U_1^{1-\rho}/(1-\rho)$. The first-order conditions are as follows:

$$(1 - \beta)c_1^{-\rho} + \frac{\beta}{1 - \alpha} [E(U_2^{1-\alpha})]^{(1-\rho)/(1-\alpha)-1} \delta'(c_1) [(\omega_0 + \omega_1 c_2^\mu)^{1-\alpha} - c_2^{1-\alpha}] = \lambda,$$

$$\beta [E(U_2^{1-\alpha})]^{(1-\rho)/(1-\alpha)-1} \{[1 - \delta(c_1)]c_2^{-\alpha} + \delta(c_1)(\omega_0 + \omega_1 c_2^\mu)^{-\alpha} \mu \omega_1 c_2^{\mu-1}\} = \lambda,$$

where β is the discount factor, α is the coefficient of relative risk aversion for static gambles, and ρ is the inverse of the elasticity of intertemporal substitution (EIS) with respect to deterministic income changes. The case of $\rho = \alpha$ and $z = 0$ corresponds to standard time-separable expected discounted utility.

In the absence of death ($\Gamma_0 = \Gamma_1 = 0$), the optimal value of the ratio c_2/c_1 is

$$\frac{c_2}{c_1} = \left(\frac{\beta}{1 - \beta} \right)^{1/\rho}. \quad (5)$$

As ρ goes to infinity (zero EIS), c_2 converges to c_1 . Suppose that $\beta > 0.5$, so people place a larger weight on the future than on the present. When ρ goes to zero (infinite EIS), c_1 converges to zero and c_2 to y , that is, all consumption takes place in period 2.

For positive values of Γ_0 and Γ_1 , the model has no analytical solution. We explore the key mechanisms using a series of numerical examples. We choose parameters so that $c_2 > \omega_0 + \omega_1 b^\mu$. This condition, emphasized by Bommier, Harenberg, and Le Grand (2020, 2021), implies that people prefer to live rather than die in the second period of their lives.

The benchmark parameters in our example are as follows: $\rho = 1/1.5$, $\alpha = 2$, $\mu = 1 - \rho$, $\omega_0 = 0.0865$, $\omega_1 = 0.1276$, $\Gamma_0 = 0.02$, $\Gamma_1 = 0.5462$. We normalize the initial income, y , to 1. Since period 2 represents the

future, we choose $\beta = 0.6$ so that more consumption occurs in the future than in the present. Given our choices of β and ρ , c_2/c_1 is equal to 1.8.

The benchmark values of ρ , α , and μ are the same as in our quantitative model. The rationale for these values is discussed in section VI.A. We choose ω_0 and ω_1 so that the following ratios coincide with the corresponding values for a weighted average of recovered young and old people in the estimated benchmark model:

$$\frac{\omega_0}{\omega_0 + \omega_1 b^\mu} = 0.44, \frac{c_2}{\omega_0 + \omega_1 b^\mu} = 3.38. \quad (6)$$

To illustrate the impact of the probability of dying on consumption, we choose values of Γ_0 and Γ_1 that are sufficiently large that the results of our experiment are clearly visible in figure 5. In our simple example, the probability of dying—evaluated at the optimum level of consumption—is quite high (20%), as is the endogenous component ($\Gamma_1 c_1$) of the probability of dying (90%). In our quantitative model, both of these numbers are much lower.

Figure 5 shows the effects of varying the EIS and risk aversion in the simple model. For each value of ρ and α , we recompute the values of ω_0 and ω_1 so that conditions (6) hold. The $\delta = 0$ line corresponds to the case in which the probability of dying in period 2 is zero ($\Gamma_0 = \Gamma_1 = 0$). The $\delta > 0$ line corresponds to the case where the probability of dying is positive and a function of c_1 ($\Gamma_0, \Gamma_1 > 0$).

The top left panel of figure 5 shows how c_1 varies with the inverse of the EIS, ρ . The dotted vertical line corresponds to the ρ value in our benchmark calibration. In general, there are two forces at work governing the impact of ρ on c_1 . First, if a person dies, they leave a bequest equal to their planned period 2 consumption. The utility of leaving this bequest is lower than that of consuming in period 2. When the EIS is high, a person reacts to the risk of dying in period 2 by reducing planned c_2 and increasing c_1 . The higher the EIS is, the larger this effect is. Second, because δ is endogenous, people have an incentive to cut c_1 to reduce the probability of dying in period 2. In our example, the second effect dominates the first effect so that for all values of ρ , c_1 is lower than when $\delta = 0$, that is, the line for the positive death case is below that of the no death case.

The top right panel of figure 5 shows how c_1 varies with the coefficient of relative risk aversion, α . When $\delta = 0$, there is no risk, so c_1 does not depend on α (the no death line is flat). When $\Gamma_0, \Gamma_1 > 0$, there are two forces governing the impact of α on c_1 . First, people respond to the risk of death by raising c_1 relative to the $\delta = 0$ case. The higher risk aversion is, the higher c_1 is. The reason is that deferring consumption to period 2 is a risky gamble relative to consuming in period 1. This effect is emphasized in Bommier, Harenberg, and Le Grand (2020). Second, in our

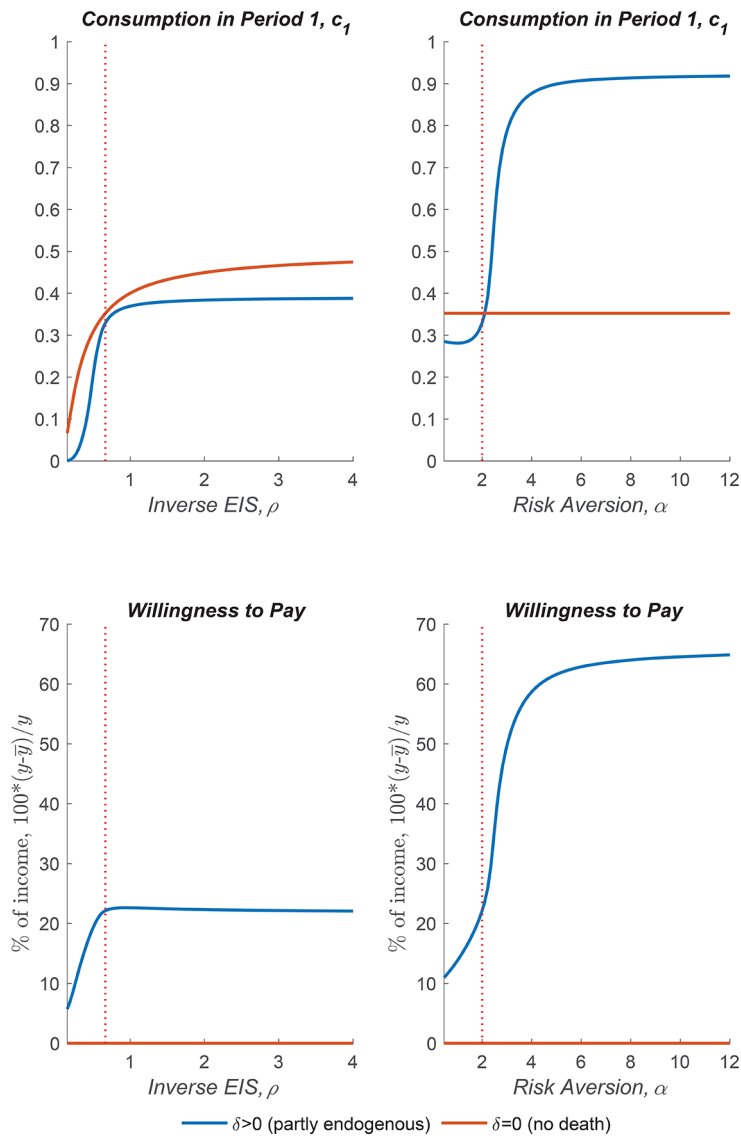


FIG. 5.—Two-period example.

model, people have an incentive to lower c_1 to reduce δ . For moderate degrees of risk aversion, the second effect dominates, so c_1 is lower than when $\delta = 0$. As risk aversion gets larger, the first effect dominates, so c_1 is higher than when $\delta = 0$.

We now turn to the question of how risk aversion and the EIS affect people's willingness to pay to eliminate the risk of death. To compute the willingness to pay, we solve the following equation, $U_1(y) = \bar{U}_1(\bar{y})$, where \bar{U}_1 is lifetime utility in an economy with $\delta = 0$. The level of \bar{y} that solves this equation is

$$\bar{y} = \frac{U_1(y)}{(1 - \beta)^{1/(1-\rho)} (1 + [\beta/(1 - \beta)]^{1/\rho})^{\rho/(1-\rho)}}.$$

The bottom left panel of figure 5 displays the fraction of income $(y - \bar{y})/y$ that people would be willing to pay to eliminate the risk of death as a function of the EIS. People's willingness to pay is low when the EIS is high (low value of ρ) because it is less costly to reduce the probability of death by cutting c_1 . For values of ρ exceeding 1, the willingness to pay is insensitive to the EIS.

The bottom right panel of figure 5 reports the analog results as we vary the coefficient of relative risk aversion, α . Not surprisingly, the willingness to pay is monotonically increasing in α . The key result is that the willingness to pay to avoid the risk of death is much more sensitive to α than to ρ . As we vary ρ , the willingness to pay ranges from roughly 8% to 22%. In contrast, as we vary α , the willingness to pay ranges from roughly 11% to 65%.

In sum, the previous discussion highlights the key mechanisms at work in our quantitative model: risk aversion, intertemporal substitution, people's beliefs about the probability of dying, and bequest motives.

VI. A Model of Consumer Behavior in an Epidemic

In this section, we develop a quantitative model of how people changed their consumption behavior in response to COVID. We use the model to address the question: how much would people be willing to pay to avoid the risk of death associated with COVID? Answering this question revolves around two issues. The first is people's beliefs about case fatality rates. The second is the fraction of the drop in consumption due to people's risk avoidance behavior as opposed to government-imposed containment measures.

We use a partial equilibrium approach that allows us to confront people of different ages and health statuses with real wages, real interest rates, and infection probabilities that mimic those observed in the data using a minimal set of assumptions. By partial equilibrium analysis, we mean that we study the consumption decisions of people of different ages and incomes, given exogenous processes for real wages, real interest rates, and infection probabilities. In section VIII, we consider a general equilibrium model in which we fully specify the environment (preferences, technology, market

structure, and epidemic dynamics) and solve for the equilibrium values of real wages, real interest rates, and infection probabilities.

Consistent with the evidence in Sorensen et al. (2022), we assume that actual case fatality rates fall over time because of improvements in medical treatments (see sec. VI.B for details). Throughout, we assume that people know the objective probability of becoming infected. However, they do not know their age group's actual time-varying case fatality rate. They begin with a prior, which they update over time. This prior and the rate at which it converges to the objective probability play a critical role in our analysis. We could have assumed that people also do not know the objective probability of becoming infected. But we could not credibly identify all the free parameters associated with this specification. As it turns out, focusing on uncertainty about the true case fatality rate is sufficient to allow the model to account for the key features of the data.

To compute the probability of being infected, people need to form expectations about the path of infections in the economy. We assume that the economy is in the preepidemic steady state in the first 4 weeks of March 2020. Then, in the fifth week of March, people learn about the first wave of the epidemic. To simplify, we assume that people have perfect foresight with respect to the first wave of infections and expect the epidemic to end in week 17 (the week of June 21, 2020). Then, in week 18 (the week of June 28, 2022), people learn that there will be two more waves. From that point on, people have perfect foresight with respect to these waves. We could allow for uncertainty about the number of infections at the cost of making the model more complex and introducing free parameters that would be difficult to identify. To add perspective on the role played by intertemporal substitution, we also consider the case in which people know there will be three waves.

We divide the population into two groups: people younger than age 60 with no comorbidities and people older than age 60 or younger than age 60 but with comorbidities. For ease of exposition, we refer to these groups as young and old. We assume that a person in the first group joins the second group with a constant probability per period, v . This assumption makes the analysis more tractable because the model has only two types of people. With deterministic aging, we would need to keep track of 61 age cohorts (from 20 to 80 years old). The critical difference between people in the two groups is the subjective and objective risk of dying from COVID or other causes.

As in Kermack and McKendrick's (1927) SIR model, people are in one of four possible health states: susceptible (those with no immunity against the virus), infected, recovered (those who recovered from the infection and have acquired immunity against the virus), and deceased. In studying the first three waves of the epidemic, we assume that recovered people have permanent immunity. This assumption is incorrect in light

of recent mutations of the COVID virus and associated breakthrough infections. However, this possibility was not widely discussed during the first three COVID waves. So, to simplify, we assume in this section that people think that once they recover from the infection, they have permanent immunity. We relax this assumption in section VIII, in which we discuss the implications of endemic COVID.

Each period in our model represents a week. Since our empirical work relies on data for public servants, we assume that people's labor supply decisions are exogenous and the real wage rate is constant. We normalize the number of hours worked to 1. The budget constraint of a person with assets b_t who consumes c_t is

$$b_{t+1} = w + (1 + r)b_t - c_t,$$

where w is the real wage rate and r is the rate of return on assets. People differ in their health status, age, and initial assets. To simplify the notation, we omit in the budget constraint the subscripts a and h .

The probability of a susceptible person in age group a becoming infected at time t , $\tau_{a,t}$, is given by the transmission function:

$$\tau_{a,t} = \pi_1 c_{a,t}^h I_t + \pi_2 I_t, \quad (7)$$

where h denotes a person's health status and I_t is the number of infected people in the population at time t . The terms $\pi_1 c_{a,t}^h I_t$ and $\pi_2 I_t$ represent the probability of becoming infected through consumption-related and non-consumption-related activities, respectively. As in Eichenbaum, Rebelo, and Trabandt (2021), this function embodies the assumption that people meet randomly and that susceptible people can reduce their infection probability by cutting their consumption.

People are uncertain about case fatality rates. At the beginning of the epidemic, people believe that the case fatality rate for a person of age a is $\pi_{ad,0}$. They update these beliefs using a parsimonious constant gain learning algorithm:⁷

$$\pi_{ad,t} = \pi_{ad,t-1} + g_a(\pi_{ad,t}^* - \pi_{ad,t-1}).$$

Here, $\pi_{ad,t}^*$ is the true case fatality rate for people of age a at time t . The parameters $g_a \in [0, 1]$ control how quickly people update their beliefs.⁸ These beliefs converge in the long run to $\pi_{ad,t}^*$. Implicitly, this specification assumes that in every period, people see the actual ratio of COVID

⁷ See Eusepi and Preston (2011) and Evans and Honkapohja (2012) for discussions of the properties of this learning algorithm.

⁸ In principle, one could entertain more complex information structures in which people receive noisy signals about infections and deaths in each period and use those signals optimally in solving their maximization problem. For computational reasons, we abstract from these types of information structures.

deaths to infections and use it to update their beliefs. At each point in time, people expect the case fatality rate to remain constant: $\pi_{ad,t+j} = \pi_{ad,t}$.⁹

The variable δ_a denotes the time t probability that a person of age a dies of non-COVID causes. The variable $\pi_{ar,t}$ denotes the probability that a person of age a who is infected at time t recovers at time $t + 1$. The probability of exiting the infection state, $\pi_{ar,t} + \pi_{ad,t}$, is constant over time, so time variation in people's beliefs about $\pi_{ad,t}$ induces time variation in their beliefs about $\pi_{ar,t}$.

As in the simple model, people's utility has the recursive form proposed by Kreps and Porteus (1978), Weil (1989), and Epstein and Zin (1991). The lifetime utility of a person with age a and health status h at time t is

$$U_{a,t}^h = \max_{c_{a,t}, b_{t+1}} \left\{ z + \left[(1 - \beta)((1 - \mu_t)c_{a,t}^h)^{1-\rho} + \beta \left\{ E_t \left[(U_{a,t+1}^h)^{1-\alpha} \right] \right\}^{(1-\rho)/(1-\alpha)} \right]^{1/(1-\rho)} \right\}. \quad (8)$$

Here, z is a constant that influences the value of life (see Hall and Jones 2007). The expectations operator, E_t , takes into account all the stochastic elements of the environment, including the possibility of death. People take as given the sequence of aggregate infections, $\{I_t\}_{t=0}^{\infty}$. We use time variation in μ_t to model exogenous changes in consumption demand associated with government-imposed containment measures. This variable represents the consumption wedge introduced by containment measures. The higher μ_t is—that is, the more containment there is—the lower the marginal utility of consumption. In appendix B.3, we show that there is an equivalence between modeling containment as a wedge on utility and a model where containment implies that some goods cannot be consumed.

The value functions for all people depend on the value of their assets, b_t , and calendar time. This time dependence reflects deterministic time variation in μ_t , I_t , $\pi_{ad,t}$, and the person's time t belief about the case fatality rates for old and young. Recall that when solving their optimization problem at time t , people assume that future values of the case fatality rate equal their current beliefs.

The value function of a susceptible young person at time t is¹⁰

$$\begin{aligned} U_{y,t}^s(b_t) = \max_{c_{y,t}, b_{t+1}} \bigg\{ & z + \{ (1 - \beta)((1 - \mu_t)c_{y,t}^s)^{1-\rho} + \beta[(1 - \tau_{y,t})(1 - \delta_y - v)(U_{y,t+1}^s(b_{t+1}))^{(1-\alpha)} \\ & + (1 - \tau_{y,t})v(U_{o,t+1}^s(b_{t+1}))^{(1-\alpha)} + \tau_{y,t}(1 - \delta_y - v)(U_{y,t+1}^i(b_{t+1}))^{(1-\alpha)} \\ & + \tau_{y,t}v(U_{o,t+1}^i(b_{t+1}))^{(1-\alpha)} + \delta_y B(b_{t+1})^{1-\alpha}]^{(1-\rho)/(1-\alpha)} \}^{1/(1-\rho)} \bigg\}. \end{aligned}$$

⁹ This assumption implies that we are working with a version of Kreps's (1998) anticipated utility.

¹⁰ This formulation and the others in app. B involve a slight abuse of notation. The perceived value function $U_{a,t+1}^h$ is computed at time t assuming that $\pi_{ad,t+j} = \pi_{ad,t}$ for all j . The realized value function at time $t + 1$ is computed assuming that $\pi_{ad,t+1+j} = \pi_{ad,t+1}$ for all j . Our notation does not distinguish between these two types of value functions. In solving the model, we do take into account this distinction.

Recall that v is the probability of a young person becoming old. U_{yt}^i and U_{ot}^i are the value functions of a young and an old infected person, respectively. The value function reflects the possible changes in health and age status at time $t + 1$. A young susceptible person at time t can remain in that state at time $t + 1$ with probability $(1 - \tau_{y,t})(1 - \delta_y - v)$, not get infected but become old with probability $(1 - \tau_{y,t})v$, get infected and stay young with probability $\tau_{y,t}(1 - \delta_y - v)$, get infected and become old with probability $\tau_{y,t}v$, or die of non-COVID causes with probability δ_y .

The function $B(b_{t+1})$ represents the utility from leaving a bequest b_{t+1} upon death. We assume that this function takes the form

$$B(b_{t+1}) = \omega_0 + \omega_1(b_{t+1})^\mu,$$

where $\omega_0 > 0$ and $\omega_1 > 0$. The bequest motive allows the model to be consistent with two empirical observations. First, many people die with large asset holdings (see, e.g., Huggett 1996; De Nardi and Yang 2014). Second, older people's consumption expenditures are lower than younger people's. The latter pattern obtains in the model because as people age, bequests receive a higher weight in the utility function relative to consumption. People of all ages and health statuses choose their consumption and asset holdings to maximize their expected lifetime utility. We solve their optimization problem using value function iteration. In appendix B, we display the value functions for old susceptible people, young infected people, old infected people, young recovered people, and old recovered people.

We partition the parameters of our quantitative model into two sets. The first set is estimated with Bayesian methods. The second set is calibrated to micro data.

A. *Parameters of Quantitative Model: Econometric Methodology*

We estimate younger and older people's initial prior beliefs about case fatality rates ($\pi_{yd,0}$ and $\pi_{od,0}$), the gain parameters (g and g_o), and the parameter μ . The latter parameter controls the impact of containment on the marginal utility of consumption. We assume that the containment wedge μ_t is given by $\mu_t = \mu \xi_t$, where μ is a scalar and ξ_t is the time series for containment measures depicted in figure 2. The maximum value of ξ_t is normalized to 1.

We calibrate the basic reproduction number, \mathcal{R}_0 , to equal 2.5, the value preferred by the Centers for Disease Control and Prevention (CDC).¹¹ In our model, \mathcal{R}_0 is given by

¹¹ See COVID-19 Pandemic Planning Scenarios, CDC, March 19, 2021.

$$\mathcal{R}_0 = \frac{\pi_1[c_{ys}s_y + c_{os}(1 - s_y)] + \pi_2}{\pi_{yr}s_y + \pi_{or}(1 - s_y) + \pi_{yd}^*s_y + \pi_{od}^*(1 - s_y)},$$

where s_y is the preepidemic share of young people in the population and c_{ys} and c_{os} are the preepidemic levels of consumption of susceptible young and old, respectively.

We estimate κ , an auxiliary parameter that represents the average share for young and old of infections generated by consumption activities at the beginning of the epidemic:

$$\kappa = \frac{\pi_1[c_{ys}s_y + c_{os}(1 - s_y)]}{\pi_1[c_{ys}s_y + c_{os}(1 - s_y)] + \pi_2}.$$

Given the value of \mathcal{R}_0 and the estimate of κ , we solve for the implied estimates of π_1 and π_2 .

Let the vector ψ denote the time series of the response to COVID of the consumption expenditures of younger and older people in our model from March 2020 to April 2021. Let $\hat{\psi}$ denote our estimate of ψ for these two groups of people obtained using regression (1). Table 12 reports the estimated regression parameters. The results are displayed in figure 8.

Our estimation criterion focuses on the consumption response of young and old with a net wealth of €75,000. According to Costa and Farinha (2012) and the Survey of Household Financial Conditions Statistics Portugal (2017), the average net wealth of Portuguese households over the period 2013–17 is €150,000. We divide this number by 2 because there are, on average, two adults per household in Portugal.

We estimate the model's predictions for people with this level of assets for two reasons. First, we do not observe the wealth distribution for people in our sample. Second, it is computationally daunting to compute the consumption behavior of people with different wealth levels in every iteration of the estimation algorithm.

The logic of the estimation procedure is conceptually the same as in Christiano, Trabandt, and Walentin (2010). Suppose that our structural model is true. Denote the true values of the model parameters by θ_0 . Let $\psi(\theta)$ denote the mapping from values of the model parameters to the time series of the impact of COVID on the consumption expenditures of younger and older people. The vector $\psi(\theta_0)$ denotes the true value of the time series whose estimates are $\hat{\psi}$. According to standard classical asymptotic sampling theory, when the number of observations, T , is large,

$$\sqrt{T}(\hat{\psi} - \psi(\theta_0)) \overset{d}{\rightarrow} N(0, W(\theta_0)).$$

It is convenient to express the asymptotic distribution of $\hat{\psi}$ as

$$\hat{\psi} \stackrel{d}{\sim} N(\psi(\theta_0), V). \quad (9)$$

Here, V is a consistent estimate of the precision matrix $W(\theta_0)/T$. Following Christiano, Trabandt, and Walentin (2010), Christiano, Eichenbaum, and Trabandt (2016), and Fernández-Villaverde, Rubio-Ramírez, and Schorfheide (2016), we assume that V is a diagonal matrix. In our case, the diagonal elements are the variances of the percentage responses of consumption of younger and older people at each point in time, reported in column 4 of table 12.

Our analysis treats $\hat{\psi}$ as observed data. We specify priors for θ and then compute the posterior distribution for θ given $\hat{\psi}$ using Bayes's rule. This computation requires the likelihood of $\hat{\psi}$ given θ . Our asymptotically valid approximation of this likelihood is motivated by (9):

$$f(\hat{\psi}|\theta, V) = (2\pi)^{-N/2} |V|^{-1/2} \exp \left[-0.5(\hat{\psi} - \psi(\theta))' V^{-1} (\hat{\psi} - \psi(\theta)) \right]. \quad (10)$$

The value of θ that maximizes this function is an approximate maximum likelihood estimator of θ . It is approximate for two reasons. First, the central limit theorem underlying (9) only holds exactly as $T \rightarrow \infty$. Second, our proxy for V is guaranteed to be correct only for $T \rightarrow \infty$.

Treating the function f as the likelihood of $\hat{\psi}$, we find that the Bayesian posterior of θ conditional on $\hat{\psi}$ and V is

$$f(\theta|\hat{\psi}, V) = \frac{f(\hat{\psi}|\theta, V)p(\theta)}{f(\hat{\psi}|V)}. \quad (11)$$

Here, $p(\theta)$ denotes the prior distribution of θ and $f(\hat{\psi}|V)$ denotes the marginal density of $\hat{\psi}$:

$$f(\hat{\psi}|V) = \int f(\hat{\psi}|\theta, V)p(\theta)d\theta.$$

Because the denominator is not a function of θ , we can compute the mode of the posterior distribution of θ by maximizing the value of the numerator in (11). We compute the posterior distribution of the parameters using a standard Monte Carlo Markov chain (MCMC) algorithm. We evaluate the relative empirical performance of different models by comparing their implications for the marginal likelihood of $\hat{\psi}$ computed using the Laplace approximation.

We assume uniform $[0, 7/14]$ priors for $\pi_{yd,0}$ and $\pi_{od,0}$ and uniform $[0, 1]$ priors for μ , g_y , g_o , and κ . We assume that it takes on average 14 days to either die or recover from an infection, so $\pi_{yd,0} + \pi_{yr,0} = 7/14$ and $\pi_{od,0} + \pi_{or,0} = 7/14$.

B. *Parameters of Quantitative Model: Calibration*

In addition to \mathcal{R}_0 , we calibrate the following parameters: $\pi_{yd,t}^*$, $\pi_{od,t}^*$, r , α , ρ , β , δ_y , δ_o , z , ω_0 , and ω_1 . We set the actual weekly case fatality rates ($\pi_{yd,t}^*$ and $\pi_{od,t}^*$) for the week of July 26, 2020, to $7 \times 0.001/14$ and $7 \times 0.035/14$, respectively. These values correspond to the case fatality rates for the median younger (age 39.5) and older (age 64.5) person (see table 1).

Sorensen et al. (2022) estimate the population-wide time trend in the infection fatality rate from April 2020 to January 2021. These estimates imply that the infection fatality rate fell by 36% between March 2022 and April 2021. We use these estimates to compute the values of $\pi_{yd,t}^*$ and $\pi_{od,t}^*$ for periods before and after July 26, 2020. We assume that the values of $\pi_{yd,t}^*$ and $\pi_{od,t}^*$ are such that, on average, infected people recover or die in 2 weeks ($\pi_{or,t}^* + \pi_{od,t}^* = \pi_{yr,t}^* + \pi_{yd,t}^* = 7/14$). We make the same assumption for the beliefs of case fatality rates, that is, $\pi_{or,t} + \pi_{od,t} = \pi_{yr,t} + \pi_{yd,t} = 7/14$. The dashed lines in figure 7 show the resulting time series for $\pi_{yd,t}^*$ and $\pi_{od,t}^*$. See appendix B.4 for a more detailed description of how we incorporate the Sorensen et al. (2022) estimates into our calibration.

The annual real interest rate, r , is set to 1%. This value corresponds roughly to the realized real yield on 10-year Portuguese government bonds from March 2020 to April 2021.

We use the life expectancy tables produced by Statistics Portugal to calibrate non-COVID-related mortality rates for younger and older people. We obtain $\delta_y = 1/(51 \times 52)$ and $\delta_o = 1/(13 \times 52)$. Since the average age difference between old and young people is 28 years, we set the weekly probability of aging, ν , to $1/(28 \times 52)$. Consistent with Portuguese demographic data, we assume that the population between 20 and 59 years old is 70% of the population between 20 and 79 years old.

We set the coefficient of relative risk aversion (α) to 2 and the EIS ($1/\rho$) to 1.5. These parameter values correspond to the estimates in Albuquerque et al. (2016), obtained using data on the equity premium and other moments of financial market data. These data are particularly relevant to our analysis because they reflect people's attitudes toward risk. The weekly discount factor, β , is set equal to $0.97^{1/52}$, which is consistent with the values used in the literature on dynamic stochastic general equilibrium models (see, e.g., Christiano, Eichenbaum, and Evans 2005).

The level parameter in the utility function (z) and the two parameters that control the utility of bequests (ω_0 and ω_1) are chosen so that the model is consistent with three features of the Portuguese data. First, the ratio of younger to older people's consumption is roughly 1.2. Second, the average savings rate is 6.7%. Third, the value of life is about €900,000, which is consistent with the value used in cost-benefit analyses of Portuguese public works (see, e.g., Ernst and Young 2015). These conditions imply

that $\omega_0 = 159.51$, $\omega_1 = 4.88$, and $z = 2.66$. A value of life of €900,000 equals 6.8 times annual consumption. For comparison, Hall, Jones, and Klenow (2020) consider values of life measured in units of consumption ranging from 5 to 7.

In our sample, the average after-tax income of people younger and older than age 60 in 2018 is very similar (€18,900 and €19,400, respectively). To simplify, we assume that both groups earn €19,000 per year.

VII. Empirical Results

Figure 6 depicts priors and posteriors for the parameters we estimate. The figure shows that the data are very informative relative to our priors. Table 2 reports the mean and 95% probability intervals for the priors and posterior of the estimated parameters.

Several features are worth noting. First, the posterior modes of $\pi_{y d, 0}$ and $\pi_{o d, 0}$ are 0.089 and 0.428, respectively. Recall that case fatality rates for young and old are $\pi_{y d} = 7 \times 0.001 / 14 = 0.0005$ and $\pi_{o d} = 7 \times 0.035 / 14 = 0.0175$, respectively. So, according to the model, both younger and

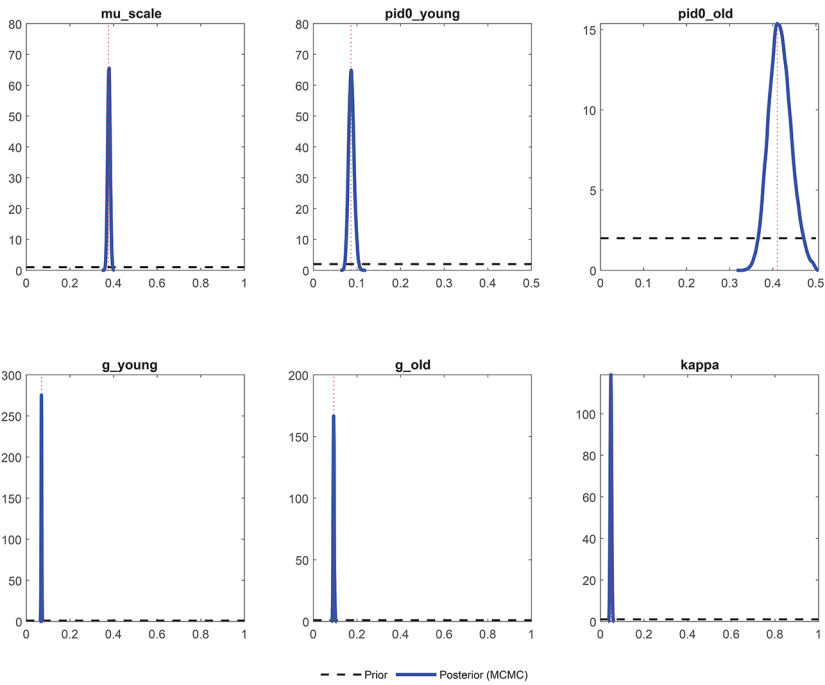


FIG. 6.—Priors and posteriors of estimated parameters.

TABLE 2
PRIORS AND POSTERiors OF PARAMETERS: BASELINE MODEL VERSUS
FIRE/NO LEARNING MODEL

	POSTERIOR DISTRIBUTION					
	PRIOR DISTRIBUTION		Baseline Model		FIRE/No Learning Model	
	Mean (1)	2.5%–97.5% (2)	Mode (3)	2.5%–97.5% (4)	Mode (5)	2.5%–97.5% (6)
Initial belief, mortality rate:						
Young, $\pi_{y,d,0}$.50	.025–.975	.087	.075–.100
Old, $\pi_{od,0}$.50	.025–.975	.413	.371–.474
Learning speed parameter:						
Young, g_y	.50	.025–.975	.069	.066–.072
Old, g_o	.50	.025–.975	.092	.088–.098
Initial share of consumption-based infections, κ	.50	.025–.975	.046	.041–.054	.069	.065–.074
Containment parameter, μ	.50	.025–.975	.380	.366–.391	.525	.518–.532
Memo item: log marginal likelihood (Laplace)				–532.5		–1,704.9

NOTE.—Ellipses indicate model specifications where particular parameter values are not relevant. Posterior mode and parameter distributions are based on a standard MCMC algorithm with a total of 500,000 draws (10 chains, 10% of draws used for burn-in, draw acceptance rates about 0.2). All prior distributions are uniform.

older people greatly overestimated their case fatality rates at the beginning of the epidemic. Second, the posterior modes of the gain parameters, g_y and g_o , are 0.069 and 0.092, respectively. Figure 7 displays the implied time series of $\pi_{y,d,t}$ and $\pi_{od,t}$. By the end of the sample, $\pi_{y,d,t}$ and $\pi_{od,t}$ have essentially converged to their true values. As discussed below, this feature is critical to the model’s ability to account for the data. Third, the posterior mode of the parameter μ is equal to 0.380. So, at their peak, containment measures reduced the marginal utility of consumption by roughly 38%. Fourth, κ , the fraction of infections associated with consumption activities is 4.6%. Taken together, these values imply that $\pi_1 = 0.000170$ and $\pi_2 = 1.1921$.

The dashed lines in figure 8 display our regression-based estimates of how the consumption of old and young people responded to COVID. The bars around point estimates represent the 95% confidence intervals. The solid lines are the corresponding model implications computed using the posterior mode of the estimated parameters. These implications are computed by comparing the model’s dynamics with and without COVID.

Figure 8 shows that the model does quite well at accounting for the consumption behavior of older people over the entire sample. In particular, the model generates the steep decline during the first wave, the recovery in the summer of 2020, the subsequent reduction beginning in

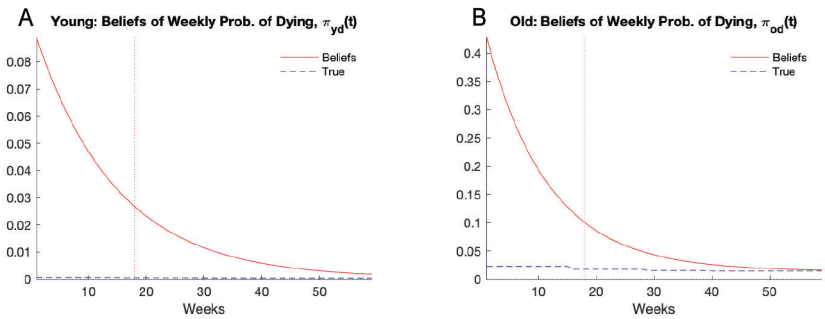


FIG. 7.—Evolution over time of beliefs about case fatality rates of old and young.

the fall of 2020, as well as the recovery in the winter of 2021. Critically, the model is consistent with the fact that consumption of the old falls by more in the first wave than in the second wave, even though the risk of infection was higher in the second and third waves.

With two exceptions, the model does quite well at accounting for the consumption behavior of the young. The first exception is that it does

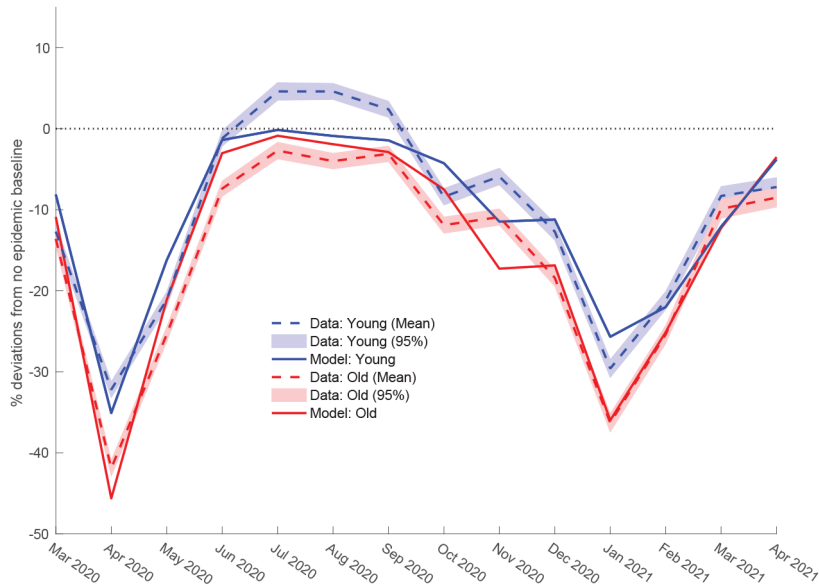


FIG. 8.—Consumption of young and old in epidemic. The figure shows the baseline estimated model and data implications for changes in expenditures of old and young during the epidemic relative to a counterfactual without COVID.

not fully explain the rise in consumption of the young during the summer of 2020. The second exception is that the model understates the peak decline in the consumption of the young during the second wave. An important success of the model is that it implies that consumption expenditures of the young fall by more in the first wave than in the second and third waves.

We conduct the following experiment to gain insight into intertemporal substitution's role in consumption choices. In our baseline specification, the second and third waves come as a surprise to people. Suppose that, instead, people knew about the second and third waves at the beginning of the epidemic. Other things equal, the more important intertemporal substitution is, the more we would expect consumption choices to be affected by this information. It turns out that in this case, people's consumption choices are very similar to the baseline case (see fig. B.8).

A. *Identifying κ and μ*

In this section, we discuss the key features of the data that allow us to identify κ and μ . Consider first κ , the fraction of infections attributed to consumption. To account for the behavior of consumers in the first and the third waves in the face of differential infections, the model assumes that initially, people have pessimistic beliefs about case fatality rates. These beliefs converge to the truth by the third wave (see fig. 7; sec. VII.B). Given this convergence, the estimation algorithm chooses κ so that the model matches the consumption of old and young in the third wave.

To understand how our model identifies μ , the parameter that controls the importance of containment, we proceed as follows. We reestimate the model, setting μ to zero. The model's fit deteriorates significantly: the marginal log likelihood falls from -532 to $-2,007$. Regarding parameter estimates, the main impact of setting μ to zero is twofold. First, it increases the value of κ , the parameter that governs the effect of consumption on the probability of being infected. Second, it reduces $\pi_{od,0}$, old people's prior about the case fatality rate.

To understand these effects, suppose that we set μ to zero without changing κ or $\pi_{od,0}$. Without containment, the consumption of the young would drop by much less than in the benchmark model. The estimation algorithm increases κ to better fit the drop in the consumption of the young. But increasing κ exacerbates the decline in the consumption of the old. To offset this effect, the estimation algorithm reduces $\pi_{od,0}$, so that the old view COVID as less lethal. The deterioration in model fit is most notable at the end of the third wave. By then, people's priors about case fatality rates have converged to their true values, and there are few infected people in the economy. In the absence of containment, consumption of young and old is counterfactually high.

We also reestimated the model by fixing μ at 10% higher than its estimated value. Even this small change in μ leads to a sizable deterioration in the marginal log likelihood, which falls from -532 to -559 . This deterioration reflects the model's poor fit at the end of the sample: consumption is counterfactually low relative to the data. We experimented with larger values of μ and found that the algorithm pushed parameters like κ and $\pi_{od,0}$ to their boundary values.

An alternative way of studying the role of containment is to compute the counterfactual fall in expenditures that would have occurred if the government had imposed containment measures, but there were no infections. The difference between the consumption policy functions with and without containment allows us to estimate the impact of containment per se. This estimate relies on the assumption that, to a first order, the observed behavior of expenditures is the sum of people's response to containment and the risk of becoming infected.

The solid line in figure 9 displays the consumption of old and young in a version of the model with containment but no infections. In this scenario, the changes in consumption expenditures of young and old people are the same. Figure 9 shows that the containment measures in isolation would have led to a 21% drop in consumption of the young and the old in the trough of the first and third waves. In the data, the actual

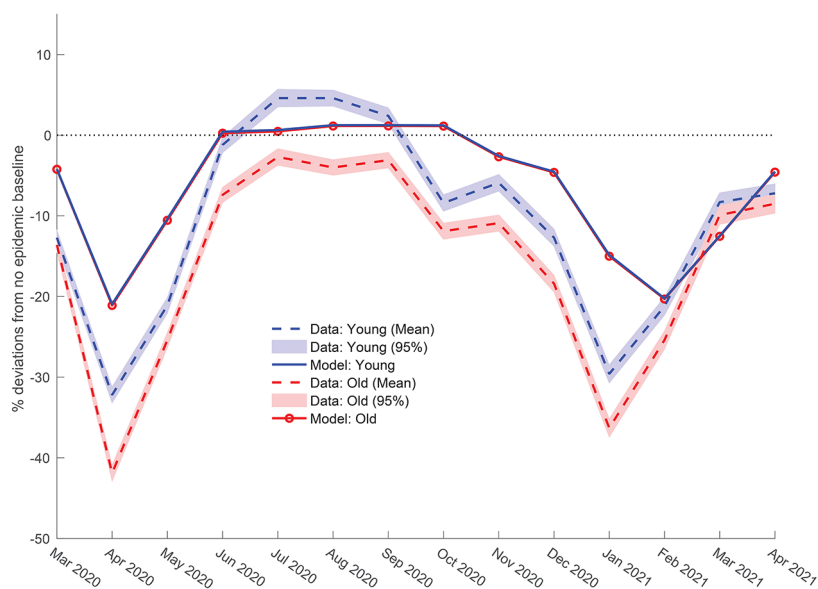


FIG. 9.—Consumption of young and old in model with containment and no COVID epidemic.

declines in consumption are much larger. So, while containment had a substantial impact, most of the decrease in consumption for both groups reflects their response to the risk of dying from COVID. These results are consistent with the findings of Arnon, Ricco, and Smetters (2020), Chetty et al. (2020), Goolsbee and Syverson (2020), Villas-Boas et al. (2020), and Chernozhukov, Kasahara, and Schrimpf (2021).

Our results are also consistent with those in Sheridan et al. (2020). Denmark and Sweden were similarly exposed to the pandemic, but only Denmark imposed significant containment measures. Sheridan et al. (2020) find that consumption of the young dropped by more in Denmark, presumably because of containment measures. Consumption of the old dropped by more in Sweden, presumably because the absence of containment increased the risk of infection.

B. The Importance of Time-Varying Beliefs

Learning plays a critical role in allowing the model to account for the key patterns in the data across the different COVID waves. In the data, the troughs of consumption are the same in the first and third waves for each age group. But the risk of becoming infected is much larger in the third wave. Other things equal, a model in which people know their true case fatality rate at the beginning of the epidemic cannot account for these facts.

To formally substantiate this claim, we estimate a version of the model with full-information rational expectations (FIRE model). In this version of the model, people know the true case fatality rates at the beginning of the epidemic. This assumption is standard in the COVID literature (e.g., Alvarez, Argente, and Lippi 2021; Eichenbaum, Rebelo, and Trabandt 2021; Jones, Philippon, and Venkateswaran 2021).

In this version of the model, the only estimated parameters are μ and κ . Columns 5 and 6 of table 2 report the mean and 95% probability intervals for the prior and posterior of μ and κ . Interestingly, the posterior modes of μ and κ are higher than the corresponding values in the benchmark model. These higher values improve model fit during the first wave but do not help the model explain the differential response of old and young.

We evaluate the performance of this model relative to the learning model by computing its implications for the marginal log likelihood. The marginal log likelihood of the no-learning model is a dramatic 1,173 points lower than that of the learning model. To understand this result, consider figure 10, which displays the implications of the reestimated model with no learning for the consumption expenditures of young and old. First, the model substantially understates the drop in consumption expenditures of old people during the first wave of the epidemic. Second, for the period up to November 2020, the model does

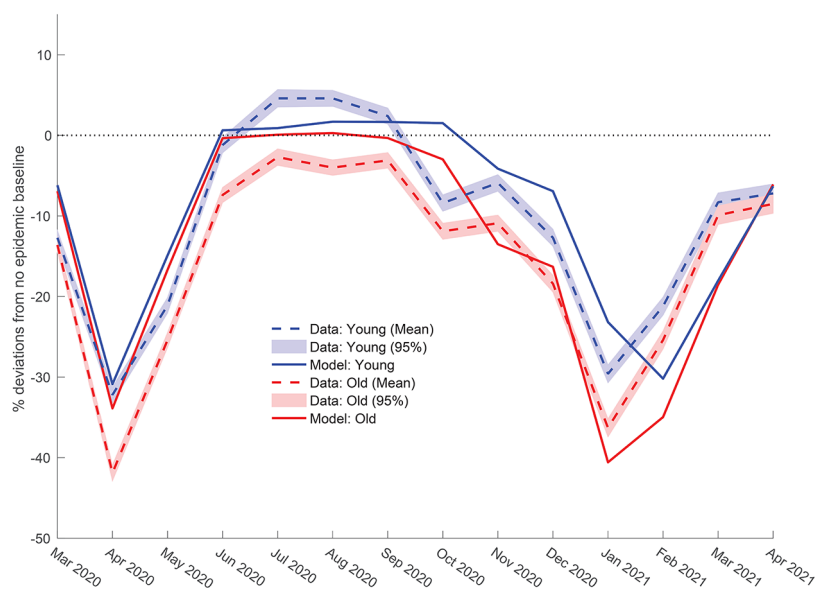


FIG. 10.—Consumption of young and old in epidemic. The figure shows the model with FIRE/no learning and data implications for changes in expenditures of young and old during the epidemic relative to a counterfactual without COVID.

not account for the fact that consumption expenditures of the old dropped by much more than those of the young. After that period, the model does generate a larger consumption drop for the old compared with the young. Third, the model counterfactually predicts that the decline in consumption expenditures of the old is larger in the second and third waves than in the first wave.

C. Consumption Response for Different Income Groups

In section IV.C, we discuss our estimates of the consumption response of different income groups to the epidemic. Recall that consumption expenditures fell more for higher-income groups than for lower-income groups. Our model is qualitatively consistent with this response pattern because higher-income households have a higher value of life, so they have more to lose from dying from COVID. In this section, we compare the quantitative implications of our model with our empirical estimates. To do so, we change the value of real labor income, w , to be consistent with the mean income of each of the three groups considered in section IV.C (€12,481, €28,566, and €59,419). We solve and simulate the model for these three income groups, keeping all parameters equal to our baseline estimates.

Figure 11 shows the model implications and the 95% confidence intervals estimated in section IV.C. This figure provides an important postestimation check on the model because these data are not used in the estimation. Except for the lower-income group during the first wave, the model fits quite well the consumption behavior of the different income groups. Introducing a subsistence level of consumption and targeted transfers to the lowest-income group would help the model better fit the consumption behavior of this group during the first wave.

D. The Alpha Variant

As a robustness check, we study the impact of alpha, the only important variant of the ancestral virus in our sample. That variant is estimated to be roughly 50% more contagious than the original strain (e.g., Yang and Shaman 2021; Tabatabai et al. 2023). Brainard et al. (2022) estimate that the alpha variant and the ancestral virus case fatality rates are roughly the same.

According to data from GISAID, this variant was detected in Portugal in the week of December 7, 2020, and consistently accounted for more

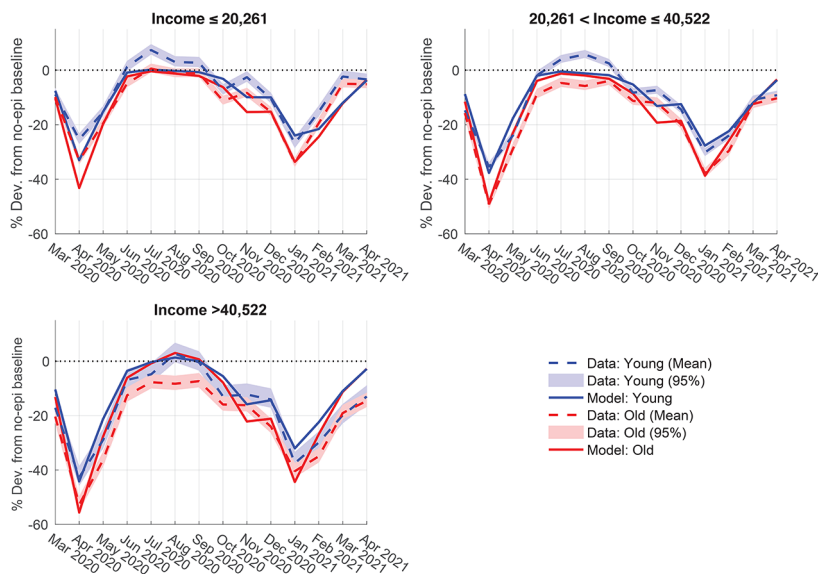


FIG. 11.—Consumption of young and old in epidemic by income groups. The figure shows the model with different levels of income and data implications for changes in expenditures of young and old with different incomes during the epidemic relative to a counterfactual without COVID.

than half of the sequenced viruses between February 2020 and April 2021.

To compute an upper bound on how much this variant affected consumption expenditures, we assume that there was an unanticipated increase of 50% in π_1 and π_2 after December 7, 2020. Figure B.9 shows that consumption by old and young falls by more in the second and third waves than in our benchmark model. This modification improves the model's fit in terms of the young's consumption expenditures and somewhat deteriorates the fit in terms of the old's consumption expenditures. Incorporating the alpha variant into the analysis does not affect our results concerning the importance of learning about case fatality rates.

E. The Impact of Declining Case Fatality Rates

In this section, we study the impact of the decline in case fatality rates estimated in Sorensen et al. (2022), which are embedded in our benchmark model. Recall that these estimates imply that the case fatality rate falls 36% between March 2020 and January 1, 2021. We solve a version of the model where the case fatality rate is constant and equal to the March 2020 values, keeping all other model parameters at their estimated baseline values. As in the baseline model, people learn the constant true case fatality rate over time. Figure B.10 displays the implications of this version of the model. The fit to the data is similar in the first wave but somewhat worse in the second and third waves. The higher case fatality rate during the second and third waves generates slightly larger consumption drops than in the baseline model. Overall, the decline in case fatality rates has a modest impact on the consumption dynamics implied by the model. The intuition for this result is that while case fatality rates declined, they did so from very low levels.

F. Willingness to Pay to Avoid the Epidemic

In this section, we study the following question: how much would people of different ages and incomes be willing to pay to avoid the epidemic? In what follows, we refer to an epidemic as including associated containment measures.

We first discuss the impact of age on the willingness to pay. The lifetime utility of a susceptible person with assets b at the beginning of the epidemic is $U_a^s(b)$. The lifetime utility of a person with assets b in an economy without an epidemic is $U_a(b)$. In general, $U_a^s(b) < U_a(b)$, that is, the epidemic reduces lifetime utility. We compute the value of initial assets \bar{b} in the economy without an epidemic that makes people indifferent between living in an economy with and without an epidemic: $U_a^s(\bar{b}) = U_a(\bar{b})$.

The annual income of a person of age a at time zero, $y_{a,0}$, is given by $y_{a,0} = 52 \times (w + rb)$. In the spirit of Murphy and Topel (2006) and Hall,

TABLE 3
WILLINGNESS TO PAY TO AVOID EPIDEMIC

$100 \times \Delta_a$	BASELINE MODEL		FIRE/NO LEARNING MODEL	
	Epidemic and Containment	No Containment	Epidemic and Containment	No Containment
Young	45	40	9	1
Old	80	77	15	7
Weighted average	54	51	10	2

Jones, and Klenow (2020), we report for young and old $\Delta \equiv (b - \bar{b})/y$, that is, the fraction of 1 year’s income that a person would be willing to give up to avoid the epidemic.

Table 3 contains our results. In the baseline model, young and old people are willing to give up 45% and 80% of a year’s income to avoid the epidemic. These large values reflect the pessimistic priors implied by the consumption behavior observed in the data. Containment increases people’s willingness to pay to avoid the epidemic, but this effect is relatively small in the baseline model. Removing containment reduces the willingness to pay from 45% to 40% for young and from 80% to 77% for old.

When people know their actual case fatality rate (FIRE/no learning), their willingness to pay to avoid the epidemic is significantly reduced to 9% and 15% of a year’s income for young and old, respectively. Since the true case fatality rate is very low for the young, most of their willingness to pay to avoid the epidemic reflects their desire to avoid the containment measures associated with the epidemic. For the old, roughly half of the willingness to pay reflects the desire to avoid containment measures.¹²

In their analysis for the United States, Hall, Jones, and Klenow (2020) compute the willingness of the representative person to pay to avoid the epidemic in a model with FIRE, no learning, and no containment. Depending on the case fatality rate, they find that a representative person would be willing to pay between 28% and 41% of 1 year’s income to avoid the epidemic. The comparable value implied by our model is 2%. Several differences between our model and that of Hall, Jones, and Klenow (2020) affect people’s willingness to pay. The two major differences are as follows.¹³ First, Hall, Jones, and Klenow (2020) assume that there is no bequest motive, so dying results in a much larger utility loss in their model. Second, they assume that the probability of dying increases by

¹² Our results on the difference in the willingness to pay of young and old are complementary to the estimates of the value of a statistical life produced by Greenberg et al. (2021) for young people in the United States who enlist in the army.

¹³ Other differences, less important from a quantitative point of view, are as follows. First, the statistical value of life and income are lower in Portugal than in the United States. Second, Hall, Jones, and Klenow (2020) assume no discounting of future utility ($\beta = 1$) and time-separable expected discounted utility.

0.81 of a percentage point for 1 year because of COVID. Their calculation corresponds to a scenario in which everybody is infected at the beginning of the epidemic and dies from COVID with probability 0.81. In our model, only a relatively small fraction of the population is infected and is at risk of dying from COVID. In our sample, the probability of dying from COVID for the overall population is 0.17%.

To illustrate the importance of two of these factors, we proceed as follows. First, we consider a version of the model with no bequests ($\omega_0 = \omega_1 = 0$) and no containment ($\mu_t = 0$). We find that the average willingness to pay is 10% of income (6% for the young and 18% for the old). Second, we assume that the number of infected is five times larger, so the probability of dying from COVID in the first year is 0.81%. We find that the average willingness to pay is 43% (29% for the young and 74% for the old), a number in the range of those reported by Hall, Jones, and Klenow (2020).

Table 4 shows how much people in different income groups would pay to avoid the epidemic. Three results emerge. First, for all income levels, the young are willing to pay less than the old both in absolute terms and as a fraction of their income. This result reflects the fact that the young are less likely to die than the old due to the epidemic. Second, the higher a person’s income is, the more they are willing to pay in absolute terms to avoid the epidemic. This finding reflects the fact that the value of life is increasing in income. Third, the higher a person’s income, the lower the fraction of their income they are willing to pay. This result reflects the fact that preferences are nonhomothetic because of the presence of two terms in the utility function (z and ω_0) that do not depend on income. These terms imply that the value of life as a fraction of income is a decreasing function of income.

VIII. The Economic Impact of Endemic COVID

In this section, we explore one way to reconcile the large short-run and small long-run effects on consumption of changes in mortality rates associated

TABLE 4
WILLINGNESS TO PAY TO AVOID EPIDEMIC BY INCOME (as Fraction of Initial Assets)

	INCOME (€1,000s)			
	12,481	19,000	28,566	59,490
€1,000s:				
Young	8	9	11	16
Old	14	16	18	25
Fraction of own income:				
Young	58	45	36	26
Old	110	80	62	41

with contagious diseases. To do so, we investigate the economic costs of endemic COVID in an economy where people know the actual case fatality rates. We modify our model in three ways. First, we modify our epidemiology model so that COVID becomes endemic. As in Abel and Panageas (2020) and Eichenbaum, Rebelo, and Trabandt (2022b), we modify social dynamics so that recovered people become susceptible with probability π_r . This modification implies that the pool of susceptible people gets replenished, so there are always new people who can get infected. As a result, the steady-state number of infected people is positive, that is, COVID is endemic. Second, we allow for vaccination. Third, we assume, for tractability, that people are organized into households, each with a continuum of identical members. This household structure introduces limited sharing of health risks. Fourth, we embed that model in a general equilibrium framework with endogenous labor choice and capital accumulation. The model is described in detail in appendix C.

Our analysis focuses on the economy’s steady state, where it seems natural to assume that people’s posteriors about case fatality rates have converged to their actual values. As might be anticipated from our previous results, this assumption has a major impact on the model’s implications for the economic consequences of endemic COVID. We compare the economic costs of COVID in this model with a counterfactual in which people have high prior values for $\pi_{yd,0}$ and $\pi_{od,0}$ and do not update them.

Steady-state results.—Column 1 of table 5 compares consumption and hours worked in the preepidemic steady state with the steady state in which COVID is endemic. Aggregate output, hours worked, and consumption fall by about 0.26% relative to the preepidemic steady state. Consumption falls by 0.86% for old people and barely falls for young people. Hours worked fall by 0.72% for older people and only 0.07% for younger people.

TABLE 5
EFFECT OF CASE FATALITY RATE ON PERCENT CHANGE OF ALLOCATIONS IN ENDEMIC
COVID STEADY STATE RELATIVE TO PREEPIDEMIC STEADY STATE

PERCENT CHANGE OF	CASE FATALITY RATE EQUAL TO	
	Actual Long-Run Value (1)	Initial Estimated Beliefs (2)
Aggregate output	−.26	−4.93
Aggregate consumption	−.26	−4.93
Aggregate hours worked	−.26	−4.93
Consumption young	−.003	−4.16
Consumption old	−.86	−6.74
Hours worked young	−.07	−4.20
Hours worked old	−.72	−6.65
Capital stock	−.26	−4.93

The lower response of consumption relative to the partial equilibrium model discussed in section VII reflects four factors. First, people know the true case fatality rate. Second, consistent with the estimates in Sorensen et al. (2022), this rate is 36% lower than at the beginning of the epidemic. Third, this model includes vaccines that reduce the probability of infection. Fourth, there are no containment measures.

We interpret these results as an upper bound on the economic cost of endemic COVID. The reason is that our model abstracts from ways in which economies can adapt to COVID. Examples include the adoption of remote work and e-commerce (for discussions, see Krueger, Uhlig, and Xie 2020; Jones, Philippon, and Venkateswaran 2021).

The steady-state economic impact of endemic COVID is minimal compared with the massive decline in economic activity experienced in 2020. In the steady state, COVID reduces life expectancy at birth by 1.5% and reduces aggregate output by 0.26% relative to the preepidemic steady state.

Interpreted through the lens of our model, the differential short- and long-run impact of endemic COVID on economic activity reflects people's beliefs about case fatality rates. The steady-state calculations above assume that people's beliefs correspond to the objective case fatality rate. Our empirical results indicate that in early 2020, people's initial beliefs about case fatality rates were much higher than the true case fatality rates.

To quantify the impact of people's beliefs on economic activity, we resolve for the steady state assuming that people make decisions on the basis of our estimates of their March 2020 beliefs. The objective case fatality rates drive actual population dynamics. Technically, in the first-order conditions for $i_{a,t+1}$, the values of π_{ad} and π_{ar} are set to the estimated initial beliefs in section VI.A.

Column 2 of table 5 compares consumption and hours worked in this steady state and the preepidemic steady state. We see large falls in consumption and hours worked relative to the preepidemic steady state. Aggregate consumption, hours worked, and physical capital fall by 4.93%. Consumption falls by 6.74% for old people and 4.2% for young people. Hours worked fall by 6.65% for older people and only 4.2% for younger people.

Taken together, our results suggest a way of reconciling the large economic impact of COVID relative to the historical evidence presented by Acemoglu and Johnson (2007). Our reconciliation highlights the critical role of expectations about case fatality rates in determining the dynamic economic impact of an epidemic.

IX. Conclusion

Our analysis highlights the importance of expectations in determining the economic impact of infectious diseases like COVID. According to our estimates, people's prior beliefs about COVID case fatality rates were

very pessimistic. These pessimistic prior beliefs led to sizable consumption declines in the first wave of the epidemic. People's beliefs converged to the true case fatality rates by the third wave of the epidemic. So, even though the risk of becoming infected was much larger in the third wave, consumption expenditures fell by about the same as in the first wave.

The fact that estimated expectations converged is important for thinking about the economic consequences of the secular declines in the mortality rate associated with infectious diseases. We expect people to eventually learn about these declines and adjust their behavior accordingly. Once this learning occurs, the impact of infectious diseases is relatively small. Our model is consistent with the large impact of COVID on economic activity and the small effect of the secular fall in mortality rates associated with other infectious diseases.

If the government and consumers have full-information rational expectations about case fatality rates, then there is a clear argument for implementing some form of containment. As discussed, for example, in Eichenbaum, Rebelo, and Trabandt (2021), there is an externality associated with market activities that can be corrected with containment measures. However, suppose that consumers overestimate case fatality rates at the beginning of an epidemic. If the government had better information than consumers, containment might be a mistake. The reason is that consumers are overreacting to the possibility of getting infected, and market activity is falling by more than warranted by the objective case fatality rate. Introducing containment could further exacerbate this overreaction.

It is unclear to us that the government had better information about case fatality rates at the beginning of the epidemic than consumers. To the extent that the government understands that it does not know the actual case fatality rates, optimal policy design should draw on the insights of the literature on decision-making under Knightian uncertainty (for a recent review, see, e.g., Gilboa and Schmeidler 1989; Ilut and Schneider 2022).

Data Availability

Code replicating the tables and figures in this article can be found in Eichenbaum et al. (2024) in the Harvard Dataverse, <https://doi.org/10.7910/DVN/EBVMFW>.

References

- Abel, Andrew B., and Stavros Panageas. 2020. "Social Distancing, Vaccination and the Paradoxical Optimality of an Endemic Equilibrium." Working Paper no. 27742, NBER, Cambridge, MA.
- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston. 2021. "Optimal Targeted Lockdowns in a Multigroup SIR Model." *A.E.R. Insights* 3 (4): 487–502.

- Acemoglu, Daron, and Simon Johnson. 2007. "Disease and Development: The Effect of Life Expectancy on Economic Growth." *J.P.E.* 115:925–85.
- Albuquerque, Rui, Martin Eichenbaum, Victor Xi Luo, and Sergio Rebelo. 2016. "Valuation Risk and Asset Pricing." *J. Finance* 71:2861–904.
- Alvarez, Fernando, David Argente, and Francesco Lippi. 2021. "A Simple Planning Problem for COVID-19 Lock-Down, Testing, and Tracing." *A.E.R. Insights* 3:367–82.
- Arnon, Alexander, John Ricco, and Kent Smetters. 2020. "Epidemiological and Economic Effects of Lockdown." *Brookings Papers Econ. Activity* 2020 (3): 61–108.
- Atkeson, Andrew, Karen Kopecky, and Tao Zha. 2020. "Four Stylized Facts about COVID-19." Working Paper no. 27719, NBER, Cambridge, MA.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Q.J.E.* 119:249–75.
- Bommier, Antoine, Daniel Harenberg, and François Le Grand. 2021. "Recursive Preferences and the Value of Life: A Clarification." Working paper.
- Bommier, Antoine, Daniel Harenberg, François Le Grand, and Cormac O'Dea. 2020. "Recursive Preferences, the Value of Life, and Household Finance." Working paper, Yale Univ., ETH Zurich.
- Brainard, J., C. M., Grossi Sampedro, A. Sweeting, and R. Fordham. 2022. "Was Alpha Deadlier than Wild-Type COVID? Analysis in Rural England." *Infection* 50:1171–78.
- Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michele Tertilt. 2021. "An Economic Model of the COVID-19 Epidemic: The Importance of Testing and Age-Specific Policies." Working paper, Bank Portugal.
- Carvalho, Vasco, Juan Ramón García, Stephen Hansen, et al. 2020. "Tracking the COVID-19 Crisis through the Lens of 1.4 Billion Transactions." VoxEU.org, April 27.
- Chernozhukov, Victor, Hiroyuki Kasahara, and Paul Schrimpf. 2021. "Causal Impact of Masks, Policies, Behavior on Early COVID-19 Pandemic in the US." *J. Econometrics* 220 (1): 23–62.
- Chetty, Raj, John N. Friedman, Nathaniel Hendren, Michael Stepner, and Opportunity Insights Team. 2020. "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." Working Paper no. 27431, NBER, Cambridge, MA.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *J.P.E.* 113:1–45.
- Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2016. "Unemployment and Business Cycles." *Econometrica* 84:1523–69.
- Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin. 2010. "DSGE Models for Monetary Policy Analysis." *Handbook Monetary Econ.* 3:285–367.
- Correia, Sergio. 2016. "A Feasible Estimator for Linear Models with Multi-Way Fixed Effects." Working paper.
- Costa, Sónia, and Luísa Farinha. 2012. "Inquérito à Situação Financeira das Famílias: Metodologia e Principais Resultados." *Banco Portugal Occasional Papers* 1:2012.
- De Nardi, Mariacristina, and Fang Yang. 2014. "Bequests and Heterogeneity in Retirement Wealth." *European Econ. Rev.* 72:182–96.
- Dessie, Zelalem G., and Temesgen Zewotir. 2021. "Mortality-Related Risk Factors of COVID-19: A Systematic Review and Meta-Analysis of 42 Studies and 423,117 Patients." *BMC Infectious Diseases* 21 (1): 855.

- Doerre, Achim, and Gabriele Doblhammer. 2022. "The Influence of Gender on COVID-19 Infections and Mortality in Germany: Insights from Age- and Gender-Specific Modeling of Contact Rates, Infections, and Deaths in the Early Phase of the Pandemic." *PLoS One* 17 (5): e0268119.
- Eichenbaum, Martin, Miguel Godinho de Matos, Francisco Lima, Sergio Rebelo, and Mathias Trabandt. 2024. "Replication Data for: 'Expectations, Infections, and Economic Activity.'" Harvard Dataverse, <https://doi.org/10.7910/DVN/EBVMFW>.
- Eichenbaum, Martin, Sergio Rebelo, and Carlos de Resende. 2017. "The Portuguese Crisis and the IMF." In *Background Papers on the IMF and the Crises in Greece, Ireland, and Portugal*, edited by M. J. Schwartz and S. Takagi, 363–447. Washington, DC: Internat. Monetary Fund.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt. 2021. "The Macroeconomics of Epidemics." *Rev. Financial Studies* 34:5149–87.
- . 2022a. "Epidemics in the New Keynesian Models." *J. Econ. Dynamics and Control* 140:104334.
- . 2022b. "The Macroeconomics of Testing and Quarantining." *J. Econ. Dynamics and Control* 138:104337.
- Epstein, Larry G., and Stanley E. Zin. 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *J.P.E.* 99:263–86.
- Ernst and Young. 2015. "Análise Custo-Benefício, Metro do Porto: Extensão Santo Ovídio—D. João II." São Paulo: Ernst and Young.
- Eusepi, Stefano, and Bruce Preston. 2011. "Expectations, Learning, and Business Cycle Fluctuations." *A.E.R.* 101:2844–72.
- Evans, George W., and Seppo Honkapohja. 2012. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton Univ. Press.
- Farboodi, Maryam, Gregor Jarosch, and Robert Shimer. 2021. "Internal and External Effects of Social Distancing in a Pandemic." *J. Econ. Theory* 196:1052–93.
- Faria-e-Castro, Miguel. 2021. "Fiscal Policy during a Pandemic." *J. Econ. Dynamics and Control* 125:1040–88.
- Fernández-Villaverde, Jesús, Juan Francisco Rubio-Ramírez, and Frank Schorfheide. 2016. "Solution and Estimation Methods for DSGE Models." In *Handbook of Macroeconomics*, vol. 2, edited by John B. Taylor and Harald Uhlig, 527–724. Amsterdam: Elsevier.
- Gaure, Simen. 2013. "lfe: Linear Group Fixed Effects." *R J.* 5:104–117.
- Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with Non-unique Prior." *J. Mathematical Econ.* 18:141–53.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull. 2020. "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic." Working Paper no. 27046, NBER, Cambridge, MA.
- Goolsbee, Austan, and Chad Syverson. 2020. "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020." Working Paper no. 27432, NBER, Cambridge, MA.
- Greenberg, Kyle, Michael Greenstone, Stephen P. Ryan, and Michael Yankovich. 2021. "The Heterogeneous Value of a Statistical Life: Evidence from U.S. Army Reenlistment Decisions." Working paper, Univ. Chicago.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2022. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" *A.E.R.* 112 (5): 1437–74.
- Guimaraes, Paulo, and Pedro Portugal. 2010. "A Simple Feasible Procedure to Fit Models with High-Dimensional Fixed Effects." *Stata J.* 10:628–49.

- Hall, Robert E., and Charles I. Jones. 2007. "The Value of Life and the Rise in Health Spending." *Q.J.E.* 122:39–72.
- Hall, Robert E., and Charles I. Jones, and J. Klenow. 2020. "Trading Off Consumption and COVID-19 Deaths." *Q. Rev.* 42 (1): 4211.
- Huggett, Mark. 1996. "Wealth Distribution in Life-Cycle Economies." *J. Monetary Econ.* 38:469–94.
- Ilut, Cosmin L., and Martin Schneider. 2022. "Modeling Uncertainty as Ambiguity: A Review." Working Paper no. 29915, NBER, Cambridge, MA.
- Jones, Callum, Thomas Philippon, and Venky Venkateswaran. 2021. "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home." *Rev. Financial Studies* 34:5188–223.
- Kermack, William Ogilvy, and Anderson G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proc. Royal Soc. London A* 115:700–721.
- Kreps, David M. 1998. "Anticipated Utility and Dynamic Choice." In *Frontiers of Research in Economic Theory*, edited by Donald P. Jacobs, Ehud Kalai, and Morton I. Kamien, 242–74. Cambridge: Cambridge Univ. Press.
- Kreps, David M., and Evan L. Porteus. 1978. "Temporal Resolution of Uncertainty and Dynamic Choice Theory." *Econometrica* 46 (1): 185–200.
- Krueger, Dirk, Harald Uhlig, and Taojun Xie. 2020. "Macroeconomic Dynamics and Reallocation in an Epidemic: Evaluating the 'Swedish Solution.'" *Econ. Policy*. 37 (110): 341–98.
- Murphy, Kevin M., and Robert H. Topel. 2006. "The Value of Health and Longevity." *J.P.E.* 114 (5): 871–904.
- Piguillem, Facundo, and Liyan Shi. 2022. "Optimal COVID-19 Quarantine and Testing Policies." *Econ. J.* 132 (647): 2534–62.
- Sheridan, Adam, Asger Lau Andersen, Emil Toft Hansen, and Niels Johannesen. 2020. "Social Distancing Laws Cause Only Small Losses of Economic Activity during the Covid-19 Pandemic in Scandinavia." *Proc. Nat. Acad. Sci. USA* 117 (34): 20468–73.
- Sorensen, Reed J. D., Ryan M. Barber, David M. Pigott, et al. 2022. "Variation in the COVID-19 Infection-Fatality Ratio by Age, Time, and Geography during the Pre-Vaccine Era: A Systematic Analysis." *Lancet* 399:1469–88.
- Statistics Portugal. 2017. "O Inquérito à Situação Financeira da Famílias (ISFF)." Bank Portugal.
- Tabatabai, Mohammad, Paul D. Juarez, Patricia Matthews-Juarez, et al. 2023. "An Analysis of COVID-19 Mortality during the Dominancy of Alpha, Delta, and Omicron in the USA." *J. Primary Care and Community Health* 14:21501319231170164.
- Toxvaerd, F. M. O. 2020. "Equilibrium Social Distancing." Working paper.
- Villas-Boas, Sofia B., James Sears, Miguel Villas-Boas, and Vasco Villas-Boas. 2020. "Are We Staying Home to Flatten the Curve?" Working paper, Univ. California, Berkeley.
- Weil, Philippe. 1989. "The Equity Premium Puzzle and the Risk-Free Rate Puzzle." *J. Monetary Econ.* 24:401–21.
- Yang, Wan, and Jeffrey Shaman. 2021. "Development of a Model-Inference System for Estimating Epidemiological Characteristics of SARS-CoV-2 Variants of Concern." *Nature Communications* 12:5573.