

Scarring Body and Mind: The Long-Term Belief-Scarring Effects of COVID-19

Julian Kozlowski, Laura Veldkamp and Venky Venkateswaran

Abstract

The largest economic cost of the COVID-19 pandemic could arise from changes in behavior long after the immediate health crisis is resolved. A potential source of such a long-lived change is scarring of beliefs, a persistent change in the perceived probability of an extreme, negative shock in the future. We show how to quantify the extent of such belief changes and determine their impact on future economic outcomes. We find that the long-run costs for the U.S. economy from this channel is many times higher than the estimates of the short-run losses in output. This suggests that, even if a vaccine cures everyone in a year, the COVID-19 crisis will leave its mark on the U.S. economy for many years to come.

Introduction

One of the most pressing questions of the day is the economic costs of the COVID-19 pandemic. While the virus will eventually pass, vaccines will be developed and workers will return to work, an event of this magnitude could leave lasting effects on the nature of economic activity. Economists are actively debating whether the recovery will be V-shaped, U-shaped or L-shaped.¹ Much of this discussion revolves around confidence, fear and the ability of firms and consumers to

rebound to their old investment and spending patterns. Our goal is to formalize this discussion and quantify these effects, both in the short- and long-run. To explore these conjectures about the extent to which the economy will rebound from this COVID-induced downturn, we use a standard economic and epidemiology framework, with one novel channel: a “scarring effect.” Scarring is a persistent change in beliefs about the probability of an extreme, negative shock to the economy. We use a version of Kozlowski et al. (2020), to formalize this scarring effect and quantify its long-run economic consequences, under different scenarios for the dynamics of the crisis.

We start from a simple premise: No one knows the true distribution of shocks in the economy. Consciously or not, we all estimate the distribution using past events, like an econometrician would. Tail events are those for which we have little data. Scarce data makes new tail event observations particularly informative. Therefore, tail events trigger larger belief revisions. Furthermore, because it will take many more observations of non-tail events to convince someone that the tail event really is unlikely, changes in tail risk beliefs are particularly persistent.

We have seen the scarring effect in action before. Before 2008, few people entertained the possibility of a financial crisis in the United States. Today, more than a decade after the global financial crisis, the possibility of another run on the financial sector is raised frequently, even though the system today is probably much safer (Baker et al. 2019). Likewise, businesses will make future decisions with the risk of another pandemic in mind. Observing the pandemic has taught us that the risks were greater than we thought. It is this newfound knowledge that has long-lived effects on economic choices.

To explore tail risk in a meaningful way, we need to use an estimation procedure that does not constrain the shape of the distribution’s tail. Therefore, we allow our agents to learn about the distribution of aggregate shocks non-parametrically. Each period, agents observe one more piece of data and update their estimates of the distribution. Section I shows how this process leads to long-lived responses of beliefs to transitory events, especially extreme, unlikely ones. The mathematical foundation for such persistence is the martingale property of beliefs. The logic is that once observed, the event remains in

agents' data set. Long after the direct effect of the shock has passed, the knowledge of that tail event affects beliefs and therefore, continues to restrain economic activity.

To illustrate the economic importance of these belief dynamics, Section II embeds our belief-updating tool in a macroeconomic model with an epidemiology event that erodes the value of capital. This framework is designed to link tail events like the current crisis to macro outcomes in a quantitatively plausible way and has been used—e.g., by Gourio (2012) and Kozlowski et al. (2020)—to study the 2007-09 Great Recession. It features, among other elements, bankruptcy risk and elevated capital depreciation from social distancing, which separates labor from capital. Section III describes the data we feed into the model to discipline our belief estimates. Section IV combines model and data and uses the resulting predictions to show how belief updating can generate large, persistent losses. We compare our results to those from the same economic model, but with agents who have full knowledge of the distribution, to pinpoint belief updating as the source of the persistence.

We model the economic effects of the COVID-19 crisis as a combination of productivity decline and accelerated capital obsolescence. We use the well-known SEIR (susceptible-exposed-infected-recovered) framework from the epidemiology literature to model the disease spread. But, it is the response to the disease that is the source of the adverse economic shock in our model. Our structure is capable of generating large asset price fluctuations, of the order observed at the onset of the pandemic, and provides a simple mapping from social distancing policies and other mitigation behavior to economic costs. It also allows us to connect to existing studies on tail risk in macroeconomics and finance. We present results for different scenarios, reflecting the considerable uncertainty about outcomes even in the short run. Our point is not to make a forecast of the coming year's events but that whatever you think will happen over the next year, the ultimate costs of this pandemic are much larger than your short-run calculations suggest.

In the first scenario, GDP drops by about 9% in 2020, recovers gradually but does not go back to its previous trajectory. It persistently

remains about 4% below the previous pre-COVID steady state. The discounted value of the lost output is almost 10 times the 2020 drop and belief revisions account for bulk of the losses (almost six times the short-run effect). Greater tail risk makes investing less attractive, reducing the stock of productive capital and (and therefore, labor input demand) persistently. In the second scenario, which captures a milder mitigation response to the spread of the disease, both short- and long-run economic costs are longer, but the relative importance of belief revisions remains the same.

The model also makes a number of predictions about asset prices. Interestingly, after an initial shock, credit spreads and equity valuations are predicted to roughly return to their original levels. This is because firms respond to this increase in riskiness by cutting back on debt. The effects of scarring are more clearly noticeable in options prices. In scenario 1, for example, the option-implied third moment in the risk-neutral distribution of equity returns becomes significantly more negative.

For monetary policymakers, one of the most pressing questions is how belief scarring will affect the long-run natural rate of interest, often referred to as “*r*-star.” Following the onset of COVID in the United States, interest rates declined rapidly. A significant portion of that decline is related to demand for liquidity. In order to understand how much of that decline was temporary and how much permanent—and more broadly about the interaction of liquidity and scarring—we introduce a role for liquid assets in an extension of our baseline model in Section V. When most capital is only partially pledgeable, but riskless assets are fully pledgeable, riskless assets, of course, have more value. But what we learn is that value is sensitive to tail risk. A persistent increase in perceived risk from COVID-19 depresses the long-run natural rate of interest by 67 basis points.

Our results also imply that a policy that prevents capital depreciation or obsolescence, even if it has only modest immediate effects on output, can have substantial long-run benefits, several times larger than the short-run considerations that often dominate policy discussion. Obviously, no policy can prevent people from believing that future pandemics are more likely than they originally thought,

but policy can change how the ongoing crisis affects capital returns. By changing that mapping, the costs of belief scarring can be mitigated. For example, bankruptcies can lead to destruction of specific investments and a permanent erosion in the value of capital. Interventions which prevent widespread bankruptcies can thus limit the adverse effects of the crisis on returns and yield substantial long-run benefits. While the short-run gains from limiting bankruptcies is well-understood, our analysis shows that neglecting the effect on beliefs leads one to drastically underestimate the benefits of such policies.

Of course, future governments could also invest in public health to mitigate the cost of future pandemics. The ability of such an investment to heal beliefs depends on the nature of belief changes induced by this episode. If we only updated our beliefs about the ability of a particular type of communicable diseases to disrupt economic activity, then health investments will be highly effective. However, traumatic events often leave survivors with a more general sense that unexpected, disastrous events can arise without warning. This more amorphous fear will be much harder for policy to combat.

Comparison to the Literature

There are many new studies of the impact of the COVID-19 pandemic on the U.S. economy, both model-based and empirical. Alvarez et al. (2020), Eichenbaum et al. (2020) and Farboodi et al. (2020) use simple economic frameworks to analyze the costs of the disease and the associated mitigation strategies. Leibovici et al. (2020) use an input-output structure to investigate the extent to which a shock to contact-intensive industries can propagate to the rest of the economy. Koren and Pető (2020) build a detailed theory-based measure of the reliance of U.S. businesses on human interaction. On the empirical side, Ludvigson et al. (2020) use VARs to estimate the cost of the pandemic over the next few months, while Carvalho et al. (2020) use high-frequency transaction data to track expenditure and behavior changes in real time. We add to this discussion by focusing on the long-term effects from changes in behavior that persist long after the disease is gone.

Other papers share our focus on long-run effects. Jorda et al. (2020) study rates of return on assets using a data set stretching to the 14th century, focusing on 15 major pandemics (with more than 100,000 deaths). Their evidence suggests a sustained downward pressure on interest rates, decades after the pandemic, consistent with long-lasting macroeconomic aftereffects. Reinhart and Rogoff (2009) examine long-lived effects of financial crises. Correia et al. (2020) find evidence of persistent declines in economic activity following the 1918 influenza pandemic. A few papers also use beliefs but rely on other mechanisms, such as financial frictions, for propagation. Elenev et al. (2020) and Krishnamurthy and Li (2020) propagate the shock primarily through financial balance sheet effects. In a more informal discussion, Cochrane (2020) explores whether the recovery from the COVID-shock will be V, U or L shaped. This work formalizes many of the ideas in that discussion.

Outside of economics, biologists and socio-biologists have noted long ago that epidemics change the behavior of both humans and animals. Loehle (1995) explore the social barriers to transmission in animals as a mode of defense against pathogen attack. Past disease events have effects on mating strategies, social avoidance, group size, group isolation and other behaviors for generations. Gangestad and Buss (1993) find evidence of similar behavior among human communities.

In the economics realm, a small number of uncertainty-based theories of business cycles also deliver persistent effects from other sorts of transitory shocks. In Straub and Ulbricht (2013) and Van Nieuwerburgh and Veldkamp (2006), a negative shock to output raises uncertainty, which feeds back to lower output, which in turn creates more uncertainty. To get even more persistence, Fajgelbaum et al. (2017) combine this mechanism with an irreversible investment cost, a combination which can generate multiple steady-state investment levels. These uncertainty-based explanations are difficult to embed in quantitative DSGE models and to discipline with macro and financial data.

Our belief formation process is similar to the parameter learning models by Johannes et al. (2016), Cogley and Sargent (2005) and Kozeniauskas et al. (2014) and is similar to what is advocated by

Hansen (2007). However, these papers focus on endowment economies and do not analyze the potential for persistent effects in a setting with production.² The most important difference is that our non-parametric approach allows us to incorporate beliefs about tail risk.

I. Belief Formation

Before laying out the underlying economic environment, we begin by explaining how we formalize the notion of belief scarring, the non-standard, but most crucial part of our analysis. We then embed it in an economic environment and quantify the effect of belief changes from the COVID-19 pandemic on the U.S. economy.

No one knows the true distribution of shocks to the economy. All of us—whether in our capacity as economic agents or modelers or econometricians—estimate such distributions, updating our beliefs as new data arrives. Our goal is to model this process in a reasonable and tractable fashion. The first step is to choose a particular estimation procedure. A common approach is to assume a normal or other parametric distribution and estimate its parameters. The normal distribution, with its thin tails, is unsuited to thinking about changes in tail risk. Other distributions raise obvious concerns about the sensitivity of results to the specific distributional assumption used. To minimize such concerns, we take a non-parametric approach and let the data inform the shape of the distribution.

Specifically, we employ a kernel density estimation procedure, one of most common approaches in non-parametric estimation. Essentially, it approximates the true distribution function with a smoothed version of a histogram constructed from the observed data. By using the widely used normal kernel, we impose a lot of discipline on our learning problem but also allow for considerable flexibility. We also experimented with a handful of other kernels.

Consider a shock $\tilde{\phi}_t$ whose true density g is unknown to agents in the economy. The agents do know that the shock $\tilde{\phi}_t$ is i.i.d. Their information set at time t , denoted \mathcal{I}_t , includes the history of all shocks $\tilde{\phi}_t$ observed up to and including t . They use this available data to construct an estimate \hat{g}_t of the true density g . Formally, at every

date, agents construct the following normal kernel density estimator of the pdf g

$$\hat{g}_t(\tilde{\phi}) = \frac{1}{n_t \kappa_t} \sum_{s=0}^{n_t-1} \Omega\left(\frac{\tilde{\phi} - \tilde{\phi}_{t-s}}{\kappa_t}\right) \quad (1)$$

where $\Omega(\cdot)$ is the standard normal density function, κ_t is the smoothing or bandwidth parameter and n_t is the number of available observations at date t . As new data arrives, agents add the new observation to their data set and update their estimates, generating a sequence of beliefs $\{\hat{g}_t\}$.

The key mechanism in the paper is the persistence of belief changes induced by transitory $\tilde{\phi}_t$ shocks. This stems from the martingale property of beliefs—i.e., conditional on time- t information (\mathcal{I}_t), the estimated distribution is a martingale. Thus, on average, the agent expects her future belief to be the same as her current beliefs. This property holds exactly if the bandwidth parameter κ_t is set to zero and holds with tiny numerical error in our application.³ In line with the literature on non-parametric assumption, we use the optimal bandwidth.⁴ As a result, any changes in beliefs induced by new information are expected to be approximately permanent. This property plays a central role in generating long-lived effects from transitory shocks.

II. Economic and Epidemiological Model

To gauge the magnitude of the scarring effect of the COVID-19 pandemic on long-run economic outcomes, we need to embed it in an economic model in which tail risk and belief changes can have meaningful effects. For this, a model needs two key features. First, it should have the potential for “large” shocks, that have both transitory and lasting effects. The former would include lost productivity from stay-at-home orders preventing services from reaching consumers. But for this shock to look like the extreme event it is to investors, the model must also allow for the possibility of a more persistent loss of productive capital. This loss represents the interior of the restaurant that went bankrupt, or the unused capacity of the hotel that will not fill again for many years to come. When stay-at-home orders forced consumers to work and consume differently, it persistently altered

tastes and habits, rendering some capital obsolete. One might think this is hard-wiring persistence in the model. Yet, as we will show, this loss of capital by itself has a short-lived effect and typically triggers an investment boom, as the economy rebuilds capital better suited to the new consumption normal. We explore two possible scenarios that highlight the enormous importance of preventing capital obsolescence, because of the scarring of beliefs.

The second key feature is sufficient curvature in policy functions, which serves to make economic activity sensitive to the probability of extreme large shocks. Two ingredients—namely, Epstein-Zin preferences and costly bankruptcy—combine to generate significant non-linearity in policy functions.

It is important to note that none of these ingredients guarantees persistent effects. Absent belief revisions, shocks, no matter how large, do not change the long-run trajectory of the economy. Similarly, the non-linear responses induced by preferences and debt influence the size of the economic response, but by themselves do not generate any internal propagation. They simply govern the magnitude of the impact, both in the short and long run.

To this setting, we add belief scarring. We model beliefs using the non-parametric estimation described in the previous section and show how to discipline this procedure with observable macro data, avoiding free parameters. This belief updating piece is not there to generate the right size reaction to the initial shock. Instead, belief updating adds the persistence, which considerably inflates the cost.

II.i. The Disease Environment

This block of the model serves to generate a time path for disruption to economic activity, which will then be mapped into transitory productivity shock and capital obsolescence. Of course, we could have directly created scenarios for the shocks and arrived at the same predictions. The explicit modeling of the spread of disease allows us to see how different social distancing policies map into shocks and ultimately into long-term economic costs from belief scarring. Given this motivation, we build on a very simple SEIR model, which is a discrete-time version of Atkeson (2020) or Stock (2020), who build on work in

the spirit of Kermack and McKendrick (1927). To this model, we add two ingredients: 1) a behavioral/policy rule that imposes capital idling when the infection rate increase (for example, this rule could represent optimal behavior or government policy); and 2) a higher depreciation rate of unused capital. While we normally think of capital utilization depreciating capital, this is a different circumstance where habits, technologies and norms are changing more rapidly than normal. Unused capital may be restaurants whose customers find new favorites, old conferencing technologies replaced with new online technology or office space that will be replaced with home offices. This higher depreciation rate represents a speeding up of capital obsolescence.

Disease and shutdowns

On Jan. 20, 2020, the first case of COVID was documented in the United States. Therefore, we start our model on that day, with one infected person. Because we are examining persistence mechanisms, we want to impose a clear end date to the COVID shock. Therefore, we assume that COVID-19 will be over by the end of 2020. The SEIR model predicts the evolution of the pandemic. Our policy shutdown rule, maps the infection rate series into a value for the aggregate shock to the U.S. economy in 2020. From 2021 onward, we assume that COVID-19 will be over. However, we explore scenarios where the economy may suffer other pandemics in the future.

Time is discrete and infinite. For the disease part of the model, we will count time in days, indexed by \tilde{t} . Later, to describe long-run effects, we will change the measure of time to t , which represents years. There are N agents in the economy. At date 1, the first person gets infected. Let S represent the number of people susceptible to the disease, but not currently exposed, infected, dead or recovered. At date 1, that susceptible number is $S(1) = N - 1$. Let E be the number of exposed persons and I be the number infected. We start with $E(1) = 0$ and $I(1) = 1$. Finally, D represents the number who are either recovered or dead, where $D(1) = 0$. The following four equations describe the dynamics of the disease.

$$S(\tilde{t} + 1) = S(\tilde{t}) - \tilde{\beta}_i S(\tilde{t}) I(\tilde{t}) / N \quad (2)$$

$$E(\tilde{t} + 1) = E(\tilde{t}) + \tilde{\beta}_i S(\tilde{t}) I(\tilde{t}) / N - \sigma_E E(\tilde{t}) \quad (3)$$

$$I(\tilde{t} + 1) = I(\tilde{t}) + \sigma_E E(\tilde{t}) - \gamma_I I(\tilde{t}) \tag{4}$$

$$D(\tilde{t} + 1) = D(\tilde{t}) + \gamma_I I(\tilde{t}) \tag{5}$$

The parameter γ_I is the rate at which people exit infection and become deceased or recovered. Thus, the expected duration of infection is approximately $1/\gamma_I$, and the number of contacts an infected person has with a susceptible person is $\tilde{\beta}$ times the fraction of the population that is susceptible $S(\tilde{t})/N$. The initial reproduction rate, often referred to as R_0 is therefore $\tilde{\beta}/\gamma_I$.

We put a t subscript on $\tilde{\beta}_t$ because behavior and policy can change it. When the infection rate rises, people reduce infection risk by staying home. This reduces the number of social contacts, reducing $\tilde{\beta}$. Lockdown policies also work by reducing $\tilde{\beta}$. We capture this relationship by assuming that $\tilde{\beta}$ can vary between a maximum of $\gamma_I R_0$ and a minimum of $\gamma_I R_{min}$. R_{min} is the estimated U.S. reproduction rate for regions under lockdown. Where on the spectrum the contact rate lies depends on the last 30-day change in infection rates, measured with a 15-day lag.⁵ Let ΔI_t be the difference between the average 15-29 day past infections and the average of 30-44 day infections: $\Delta I_t = (1/15) \left(\sum_{\tau=15}^{29} I(t-\tau) - \sum_{\tau=30}^{44} I(t-\tau) \right)$. This captures the fact that most policymakers are basing policy on two-week changes in hospitalization rates, which are themselves observed with a 14-day lag. Then policy and individual behavior achieves a frequency of social contact:

$$\tilde{\beta}_t = \gamma_I \times \min(R_0, \max(R_{min}, R_0 - \zeta^* \Delta I_t)) \tag{6}$$

The key part of the epidemic from a belief-scarring perspective is that reducing the contact rate requires separating labor from capital. In other words, capital is idled (full capacity) when no mitigation efforts are underway, i.e. when $\tilde{\beta}_t = \gamma_I R_0$. But as $\tilde{\beta}_t$ falls, capital idling (K^-) rises. We formalize that relationship as

$$K^-_t = \tilde{\theta}^* (R_0 - \tilde{\beta}_t / \gamma_I). \tag{7}$$

Idle capital depreciates as a rate $\tilde{\delta}$. As mentioned before, this is not physical deterioration of the capital stock. Instead, it represents a loss of value from accelerated obsolescence due to changes in tastes, habits and technologies. It could also represent a loss in value because of persistent upstream or downstream supply chain constraints.

II.ii. The Economy

Preferences and technology

To describe long-term economic consequences, we switch from the daily time index \tilde{t} to an annual time index t . An infinite horizon, discrete time economy has a representative household, with preferences over consumption (C_t) and labor supply (L_t):

$$U_t = \left[(1-\beta)(c_t^\gamma (1-L_t)^{1-\gamma})^{1-\Psi} + \beta E_t (U_{t+1}^{1-\eta})^{\frac{1-\Psi}{1-\eta}} \right]^{\frac{1}{1-\Psi}} \quad (8)$$

where Ψ is the inverse of the inter-temporal elasticity of substitution, η indexes risk-aversion, γ indexes the share of consumption in the period utility function, and β represents time preference.

The economy is also populated by a unit measure of firms, indexed by i and owned by the representative household. Firms produce output with capital and labor, according to a standard Cobb-Douglas production function $z_t k_{it}^\alpha l_{it}^{1-\alpha}$.

Aggregate uncertainty is captured by a single random variable, $\tilde{\phi}_t$, which is i.i.d over time and drawn from a distribution $g(\cdot)$. The i.i.d. assumption is made in order to avoid an additional, exogenous, source of persistence.⁶ The effect of this shock on economic activity depends on the realized default rate Def_t^f (the fraction of firms who default in t , characterized later in this section). Formally, it induces a capital obsolescence “shock” $\phi_t \equiv \Phi(\tilde{\phi}_t, Def_t^f)$. The function $\Phi(\cdot)$ will be made explicit later. This composite shock has both permanent and transitory effects. The permanent component works as follows: a firm that enters the period t with capital \hat{k}_{it} has effective capital $k_{it} = \phi_t \hat{k}_{it}$.

In addition to this permanent component, the shock ϕ_t also has a temporary effect, through the TFP term $z_t = \phi_t^\nu$. The parameter ν

governs the relative strength of the transitory component. This specification allows us to capture both permanent and transitory disruptions with only one source of uncertainty. By varying v , we can capture a range of scenarios without having to introduce additional shocks.

Firms are also subject to an idiosyncratic shock v_{it} . These shocks scale up and down the total resources available to each firm (after paying labor, but before paying debtholders' claims)

$$\Pi_{it} = v_{it} [z_t k_{it}^\alpha l_{it}^{1-\alpha} - W_t l_{it} + (1-\delta)k_{it}] \quad (9)$$

where δ is the ordinary rate of capital depreciation. The additional obsolescence from idle capital is already removed from k_{it} , via the shock ϕ_i . The shocks v_{it} are i.i.d. across time and firms and are drawn from a known distribution, F .⁷ The mean of the idiosyncratic shock is normalized to be one: $\int v_{it} di = 1$. The primary role of these shocks is to induce an interior default rate in equilibrium, allowing a more realistic calibration, particularly of credit spreads.

What is capital obsolescence?

Capital obsolescence shock reflects a long-lasting change in the economic value of the average unit of capital. A realization of $\phi < 1$ captures the loss of specific investments or other forms of lasting damage from a prolonged shutdown. This could come from the lost value of cruise ships that will never sail again, businesses that do not re-open, loss of customer capital or just less intensive use of commercial space due to a persistent preference for more distance between other diners, travelers, spectators or shoppers. It could also represent permanent changes in health and safety regulations that make transactions safer, but less efficient from an economic standpoint.

An important question is whether future investment could be made in ways or in sectors that avoid these costs. Of course, such substitution is likely to happen to some extent. But, the fact that the patterns of investment were not chosen previously suggests that these adjustments are costly or less profitable. More importantly, we learned that the world is riskier and more unpredictable than we thought. The

shocks that hit one sector (or type of capital) today may hit another tomorrow, in ways that are impossible to foresee.

Capital markets and default

Firms have access to a competitive non-contingent debt market, where lenders offer bond price (or equivalently, interest rate) schedules as a function of aggregate and idiosyncratic states, in the spirit of Eaton and Gersovitz (1981). A firm enters period $t+1$ with an obligation, b_{it+1} to bondholders. The shocks are then realized and the firm's shareholders decide whether to repay their obligations or default. Default is optimal for shareholders if and only if

$$\Pi_{it+1} - b_{it+1} + \Gamma_{t+1} < 0$$

where Γ_{t+1} is the present value of continued operations. Thus, the default decision is a function of the resources available to the firm Π_{it+1} (output plus undepreciated capital less wages) and the obligations to bondholders b_{it+1} . Let $r_{it+1} \in \{0, 1\}$ denote the default policy of the firm.

In the event of default, equity holders get nothing. The productive resources of a defaulting firm are sold to an identical new firm at a discounted price, equal to a fraction $\theta < 1$ of the value of the defaulting firm. The proceeds are distributed pro-rata among the bondholders.⁸

Let q_{it} denote the bond price schedule faced by firm i in period t , i.e., the firm receives q_{it} in exchange for a promise to pay one unit of output at date $t+1$. Debt is assumed to carry a tax advantage, which creates incentives for firms to borrow. A firm which issues debt at price q_{it} and promises to repay b_{it+1} in the following period, receives a date- t payment of $\chi q_{it} b_{it+1}$, where $\chi > 1$. This subsidy to debt issuance, along with the cost of default, introduces a trade-off in the firm's capital structure decision, breaking the Modigliani-Miller theorem.⁹

For a firm that does not default, the dividend payout is its total available resources, minus its payments to debt and labor, minus the cost of building next period's capital stock (the undepreciated current capital stock is included in Π_{it}), plus the proceeds from issuing new debt, including its tax subsidy

$$d_{it} = \Pi_{it} - b_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1}. \quad (10)$$

Importantly, we do not restrict dividends to be positive, with negative dividends interpreted as (costless) equity issuance. Thus, firms are not financially constrained, ruling out another potential source of persistence.

Bankruptcy and obsolescence

Next, we spell out the relationship between default and capital obsolescence, $\phi_t = \Phi(\tilde{\phi}_t, Def_t)$ where $Def_t \equiv \int r_{it} di$. This is meant to capture the idea that widespread bankruptcies can amplify the erosion in the economic value of capital arising from the primitive shock $\tilde{\phi}_t$. This might come from lost supply chain linkages, interfirm relationships or other ways in which economic activity is interconnected. For example, a retailer might ascribe a lower value to space in a mall if a number of other stores go out of business. Similarly, a manufacturer might need to undertake costly search or make adjustments to his factory in order to accommodate new suppliers. We capture these effects with a flexible functional form:

$$\ln \phi_t = \ln \Phi(\tilde{\phi}_t, Def_t) = \ln \tilde{\phi}_t - \mu Def_t^{1-\sigma}, \tag{11}$$

where μ and σ are parameters that govern the relationship between default and the loss of capital value.

Timing and value functions:

1. Firms enter the period with a capital stock \hat{k}_{it} and outstanding debt b_{it} .
2. The aggregate capital obsolescence shocks are realized.¹⁰ Labor choice is made and production takes place.
3. Firm-specific shocks v_{it} are realized. The firm decides whether to default or repay ($r_{it} \in \{0, 1\}$) its debt claims and distribute any remaining dividends.
4. The firm makes capital \hat{k}_{it+1} and debt b_{it+1} choices for the following period.

In recursive form, the problem of the firm is

$$V(\hat{k}_{it}, b_{it}, v_{it}, S_t) = \max \left[0, \max_{d_{it}, J_{it}, \hat{k}_{it+1}, b_{it+1}} d_{it} + \mathbb{E}_t M_{t+1} V(\hat{k}_{it+1}, b_{it+1}, v_{it+1}, S_{t+1}) \right] \tag{12}$$

where M_{t+1} is the representative households' stochastic discount factor, subject to

$$\text{Dividends: } d_{it} \leq \Pi_{it} - b_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1} \quad (13)$$

$$\text{Resources: } \Pi_{it} = v_{it} [z_t k_{it}^\alpha l_{it}^{1-\alpha} - W_t l_{it} + (1-\delta)k_{it}] \quad (14)$$

$$\text{Bond price: } q_{it} = \mathbb{E}_t M_{t+1} \left[r_{it+1} + (1-r_{it+1}) \frac{\theta \tilde{V}_{it+1}}{b_{it+1}} \right] \quad (15)$$

Finally, firms hire labor in a competitive market at a wage W_t . We assume that this decision is made after observing the aggregate shock but before the idiosyncratic shocks are observed, i.e., labor choice solves the following static problem:

$$\max_{l_{it}} z_t (\phi_t \hat{k}_{it})^\alpha l_{it}^{1-\alpha} - W_t l_{it}$$

The first max operator in (12) captures the firm's option to default. The expectation \mathbb{E}_t is taken over the idiosyncratic and aggregate shocks, given beliefs about the aggregate shock distribution. The value of a defaulting firm is simply the value of a firm with no external obligations, i.e., $\tilde{V}_{it} = V(\hat{k}_{it}, 0, v_{it}, S_t)$.

The aggregate state S_t consists of $(\hat{K}_t, \tilde{\phi}_t, \mathcal{I}_t)$ where \mathcal{I}_t is the economy-wide information set. Equation (15) reveals that bond prices are a function of the firms capital \hat{k}_{it+1} and debt b_{it+1} , as well as the aggregate state S_t . The firm takes the aggregate state and the function $q_{it} = q(\hat{k}_{it+1}, b_{it+1}, S_t)$ as given, while recognizing that its choices affect its bond price.

Information, beliefs and equilibrium

The set \mathcal{I}_t includes the history of all shocks $\tilde{\phi}$ observed up to and including time- t . The expectation operator \mathbb{E}_t is defined with respect to this information set. Expectations are probability-weighted integrals, where the probability density is $\hat{g}(\tilde{\phi})$. The function \hat{g} arises from using the kernel density estimation procedure in equation (1).

For a given belief \hat{g} , a recursive equilibrium is a set of functions for (i) aggregate consumption and labor that maximize (8) subject to a

budget constraint, (ii) firm value and policies that solve (12), taking as given the bond price function (15) and the stochastic discount factor, (iii) aggregate consumption and labor are consistent with individual choices and (iv) capital obsolescence is consistent with default rates according to (11).

II.iii. Characterization

The equilibrium of the economic model is a solution to the following set of nonlinear equations. First, the fact that the constraint on dividends (13) will bind at the optimum can be used to substitute for d_{it} in the firm's problem (12). This leaves us with two intertemporal choice variables (\hat{k}_{it+1}, b_{it+1}) and a default decision. The latter is described by a threshold rule in the idiosyncratic output shock v_{it} :

$$r_{it} = \begin{cases} 0 & \text{if } v_{it} < \underline{v} \\ 1 & \text{if } v_{it} \geq \underline{v} \end{cases}$$

which implies that the default rate $Def_t = F(\underline{v})$. It turns out to be more convenient to redefine variables and cast the problem as a choice of \hat{k}_{it+1} and leverage, $lev_{it+1} \equiv \frac{b_{it+1}}{\hat{k}_{it+1}}$. The full characterization to the Appendix. Since all firms make symmetric choices for these objects, in what follows, we suppress the i subscript. The optimality condition for \hat{k}_{t+1} is:

$$1 = \mathbb{E}[M_{t+1}R_{t+1}^k] + (\chi - 1)lev_{t+1}q_t - (1 - \theta)\mathbb{E}[M_{t+1}R_{t+1}^k h(\underline{v}_{t+1})] \quad (16)$$

$$\text{where } R_{t+1}^k = \frac{\phi_{t+1}^{\alpha+\nu} \hat{k}_{t+1}^\alpha l_{t+1}^{1-\alpha} - W_{t+1} l_{t+1} + (1 - \delta)\phi_{t+1} \hat{k}_{t+1}}{\hat{k}_{t+1}} \quad (17)$$

The object R_{t+1}^k is the ex-post per-unit, post-wage return on capital, which is obviously a function of the obsolescence shock ϕ_t . The default threshold is given by $\underline{v}_{t+1} = \frac{lev_{t+1}}{R_{t+1}^k}$ while $h(\underline{v}) \equiv \int_{-\infty}^{\underline{v}} v f(v) dv$ is the default-weighted expected value of the idiosyncratic shock.

The first term on the right hand side of (16) is the usual expected direct return from investing, weighted by the stochastic discount factor. The other two terms are related to debt. The second term reflects

the indirect benefit to investing arising from the tax advantage of debt—for each unit of capital, the firm raises $\frac{b_{t+1}}{\hat{k}_{t+1}}q_t$ from the bond market and earns a subsidy of $\chi-1$ on the proceeds. The last term is the cost of this strategy—default-related losses, equal to a fraction $1-\theta$ of available resources.

Note that the default threshold is a function of ϕ_t , which in turn is affected by default, through (11). Therefore, the threshold equation $\underline{v}_{t+1} = \frac{lev_{t+1}}{R_{t+1}^k}$ implicitly defines a fixed-point relationship:

$$\underline{v}_{t+1} = \frac{lev_{t+1}}{R_{t+1}^k} = \frac{lev_{t+1}}{\phi_{t+1}^{\alpha+v} \hat{k}_{t+1}^{\alpha-1} l_{t+1}^{1-\alpha} - W_{t+1} \frac{l_{t+1}}{\hat{k}_{t+1}} + (1-\delta)\phi_{t+1}} \quad (18)$$

Next, the firm's optimal choice of leverage, lev_{t+1} is

$$(1-\theta)\mathbb{E}_t \left[M_{t+1} \frac{lev_{t+1}}{R_{t+1}^k} f \left(\frac{lev_{t+1}}{R_{t+1}^k} \right) \right] = \left(\frac{\chi-1}{\chi} \right) \mathbb{E}_t \left[M_{t+1} \left(1 - F \left(\frac{lev_{t+1}}{R_{t+1}^k} \right) \right) \right]. \quad (19)$$

The left hand side is the marginal cost of increasing leverage—it raises the expected losses from the default penalty (a fraction $1-\theta$ of the firm's value). The right hand side is the marginal benefit—the tax advantage times the value of debt issued.

Finally, firm and household optimality implies that labor solves the intra-temporal condition:

$$\frac{(1-\alpha)y_t}{l_t} = W_t = \frac{1-\gamma}{\gamma} \frac{c_t}{1-l_t} \quad (20)$$

The optimality conditions, (16)–(20), along with those from the household side, form the system of equations we solve numerically.

III. Measurement, Calibration and Solution Method

This section describes how we use macro data to estimate beliefs and parameterize the model, as well as our computational approach. A

strength of our theory is that we can use observable data to estimate beliefs at each date.

Measuring past shocks

Of course, we have not seen a health event like COVID in the last 95-100 years. However, from an economic point of view, COVID is one of many past shocks to returns that happens to be larger. When we think about COVID changing our beliefs, or our perceived probability distribution of outcomes, those outcomes are realized returns on capital. Therefore, to estimate the pre-COVID and post-COVID probability distributions, we first set out to measure past capital returns that map neatly into our model.

A helpful feature of capital obsolescence shocks, like the ones in our model, is that their mapping to available data is straightforward. A unit of capital installed in period $t - 1$ (i.e., as part of \hat{k}_t) is, in effective terms, worth ϕ_t units of consumption goods in period t . Thus, the change in its market value from $t - 1$ to t is simply ϕ_t .

We apply this measurement strategy to annual data on commercial capital held by U.S. corporates. Specifically, we use two time series Non-residential assets from the Flow of Funds, one evaluated at market value and the second, at historical cost.¹¹ We denote the two series by NFA_t^{MV} and NFA_t^{HC} respectively. To see how these two series yield a time series for ϕ_t , note that, in line with the reasoning above, NFA_t^{MV} maps directly to effective capital in the model. Formally, letting P_t^k be the nominal price of capital goods in t , we have $P_t^k K_t = NFA_t^{MV}$. Investment X_t can be recovered from the historical series, $P_{t-1}^k X_t = NFA_t^{HC} - (1 - \delta)NFA_{t-1}^{HC}$. Combining, we can construct a series for $P_{t-1}^k \hat{K}_t$:

$$\begin{aligned} P_{t-1}^k \hat{K}_t &= (1 - \delta)P_{t-1}^k K_{t-1} + P_{t-1}^k X_t \\ &= (1 - \delta)NFA_{t-1}^{MV} + NFA_t^{HC} - (1 - \delta)NFA_{t-1}^{HC} \end{aligned}$$

Finally, in order to obtain $\phi_t = \frac{K_t}{\hat{K}_t}$, we need to control for nominal price changes. To do this, we proxy changes in P_t^k using the price index for non-residential investment from the National Income and Product Accounts (denoted PINDX).¹² This yields:

$$\begin{aligned}\phi_t &= \frac{K_t}{\dot{K}_t} = \left(\frac{P_t^k K_t}{P_{t-1}^k \dot{K}_t} \right) \left(\frac{PINDX_{t-1}^k}{PINDX_t^k} \right) \\ &= \left[\frac{NFA_t^{MV}}{(1-\delta)NFA_{t-1}^{MV} + NFA_t^{HC} - (1-\delta)NFA_{t-1}^{HC}} \right] \left(\frac{PINDX_{t-1}^k}{PINDX_t^k} \right)\end{aligned}\quad (21)$$

Using the measurement equation (21), we construct an annual time series for capital depreciation shocks for the U.S. economy since 1950. The mean and standard deviation of the series over the entire sample are 1 and 0.03, respectively. The autocorrelation is statistically insignificant at 0.15.

Next, we recover the primitive shock $\tilde{\phi}_t$ from the time series ϕ_t . To do this, we use (11), along with data on historical default rates from Moody's Investors Service (2015)¹³ and values for the feedback parameters (μ, ϖ) as described below. The first panel of Chart 2 shows the estimated $\tilde{\phi}$.

Parameterization

A period t is interpreted as a year. We choose the discount factor $\beta = 0.95$, depreciation $\delta = 0.06$, and the share of capital in the production, α , is 0.40. The recovery rate upon default, θ , is set to 0.70, following Gourio (2013). The distribution for the idiosyncratic shocks, v_{it} is assumed to be lognormal, i.e., $\ln v_{it} \sim N\left(-\frac{\hat{\sigma}^2}{2}, \hat{\sigma}^2\right)$ with $\hat{\sigma}^2$ chosen to target a default rate of 0.02.¹⁴ The share of consumption in the period utility function, γ , is set to 0.4.

For the parameters governing risk aversion and intertemporal elasticity of substitution, we use standard values from the asset pricing literature and set $\psi = 0.5$ (or equivalently, an IES of 2) and $\eta=10$. The tax advantage parameter χ is chosen to match a leverage target of 0.50, the ratio of external debt to capital in the U.S. data—from Gourio (2013). Finally, we set the parameters of the default-obsolescence feedback function, namely μ and ϖ . Ideally, these parameters would be calibrated to match the variability of default and its covariance with the observed ϕ_t shock. Unfortunately, our one-shock model fails to generate enough volatility in default rates and therefore, struggles to match these moments. Fixing this would almost certainly require

a richer model with multiple shocks and more involved financial frictions. We take a simpler way out here and target a relatively modest feedback with values of $\mu = 0.2$ and $\varpi = 0.5$. These values imply roughly an amplification 3% at a baseline default rate of 2%, rising to 5% for a 6% default.¹⁵

Epidemiology parameters

A major hurdle to quantifying the long-run effects is the lack of data and uncertainty surrounding estimates of the short-run impact. While this will surely be sorted out in the months to come, for now, with the crisis still raging and policy still being set, the impact is uncertain. More importantly for us, the nature of the economic shock is uncertain. It may be a temporary closure with furloughs, or it could involve widespread bankruptcies and changes in habits that permanently separate workers from capital or make the existing stock of capital ill-suited to the new consumption demands. Since it is too early to know this, we present two possible scenarios, chosen to illustrate the interaction between learning and the type of shock we experience. All involve significant losses in the short term but their long-term effects on the economy are drastically different.

We begin by describing parameter choices that are fixed across the scenarios. Following Wang et al. (2020)'s study of infection in Hubei, China, we calibrate $\sigma_E = 1/5.2$ and $\gamma_I = 1/18$ to the average duration of exposure (5.2 days) and infection (18 days). We use an initial reproduction number of $R_0 = 3.5$, based on more recent estimates of higher antibody prevalence and more asymptomatic infection than originally thought and $R_{min} = 0.8$ based on the estimates of the spread in New York, at the peak of the lockdown (Center for Disease Control 2020). This implies that the initial number of contacts per period must be $\tilde{\beta} = \gamma_I R_0$.

The extent to which capital idling reduces contact rates is set to $\tilde{\theta} = 1/3$. This implies that a lockdown which reduces the reproduction number to 0.8 is associated with 50% capital idling. This is broadly consistent with the 25% drop in output, estimated during the lockdown period in Hubei province, China. The rate of excess depreciation

of idle capital at the rate of 6.5% per month or $\tilde{\delta} = 0.065/30$ daily. As we will see, this implies a 10% erosion of the value of capital in our first scenario, which lines up with the drop in commercial real estate prices since the pandemic started—see CPPI (2020).

The two scenarios, which differ in the sensitivity of lockdown policy to observed infection increases, i.e., the parameter ζ_I . In scenario 1, we set $\zeta_I = 300$, which generates an initial lockdown that lasts for two months. This version of the model predicts waves of re-infection and new lockdowns in the months to come, echoing predictions by the Centers for Disease Control. Scenario 2, which considers a much less aggressive response by setting $\zeta_I = 50$, has only one lockdown episode.

Numerical solution method

Table 1 summarizes the resulting parameter choices. Since curvature in policy functions is an important feature of the economic environment, our algorithm solves equations (20)–(19) with a non-linear collocation method. Appendix A.B describes the iterative procedure. In order to keep the computation tractable, we need one more approximation. The reason is that date- t decisions (policy functions) depend on the current estimated distribution ($\hat{g}_t(\tilde{\phi})$) and the probability distribution h over next-period estimates, $\hat{g}_{t+1}(\tilde{\phi})$. Keeping track of $h(\hat{g}_{t+1}(\tilde{\phi}))$, (a compound lottery) makes a function a state variable, which renders the analysis intractable. However, the approximate martingale property of \hat{g}_t discussed in Section I offers an accurate and computationally efficient approximation to this problem. The martingale property implies that the average of the compound lottery is $E_t[\hat{g}_{t+1}(\tilde{\phi})] \approx \hat{g}_t(\tilde{\phi}), \forall \tilde{\phi}$. Therefore, when computing policy functions, we approximate the compound distribution $h(\hat{g}_{t+1}(\tilde{\phi}))$ with the simple lottery $\hat{g}_t(\tilde{\phi})$, which is today's estimate of the probability distribution. We use a numerical experiment to show that this approximation is quite accurate. The reason for the small approximation error is that $h(\hat{g}_{t+1})$ results in distributions centered around $\hat{g}_t(\tilde{\phi})$, with a small standard deviation. Even 30 periods out, $\hat{g}_{t+30}(\tilde{\phi})$ is still quite close to its mean $\hat{g}_t(\tilde{\phi})$. For 1-10 years ahead, where most of the utility weight is, this standard error is tiny.

Table 1
Parameters

Parameter	Value	Description
Preferences:		
β	0.95	Discount factor
η	10	Risk aversion
ψ	0.50	1/Intertemporal elasticity of substitution
γ	0.40	Share of consumption in the period utility function
Technology:		
α	0.40	Capital share
δ	0.06	Depreciation rate
$\hat{\sigma}$	0.28	Idiosyncratic volatility
Debt:		
χ	1.06	Tax advantage of debt
θ	0.70	Recovery rate
μ	0.2	Default-obsolescence feedback
ϖ	0.5	Default-obsolescence elasticity
Disease / Policy:		
R_0	3.5	Initial disease reproduction rate
R_{min}	0.8	Minimum U.S. disease reproduction rate
σ_E	1/52	Exposure to infection transition rate
γ_I	1/18	Recovery / death rate
ξ_I	300 (50)	Lockdown policy sensitive to past infections
$\tilde{\theta}$	0.19	Capital idling required to reduce transmission
$\tilde{\delta}$	0.002	Excess depreciation (daily) of idle capital

Notes: The number in parentheses is used in scenario 2.

To compute our benchmark results, we begin by estimating \hat{g}_{2019} using the data on $\tilde{\phi}_t$ described above. Given this \hat{g}_{2019} , we compute the stochastic steady state by simulating the model for 5,000 periods, discarding the first 500 observations and time-averaging across the remaining periods. This steady state forms the starting point for our results. Subsequent results are in log deviations from this steady state level. Then, we subject the model economy to two possible additional adverse realizations for 2020, one at a time. Using the one additional data point for each scenario, we re-estimate the distribution, to get \hat{g}_{2020} . To see how persistent economic responses are, we need

a long future time series. We don't know what distribution future shocks will be drawn from. Given all the data available to us, our best estimate is also \hat{g}_{2020} . Therefore, we simulate future paths by drawing many sequences of future $\tilde{\phi}$ shocks from the \hat{g}_{2020} distribution. In the results that follow, we plot the mean future path of various aggregate variables.

IV. Main Results

Our goal in this paper is to quantify the long-run effect of the COVID crises, stemming from the belief scarring effect, i.e., from learning that pandemics are more likely than we thought. We formalize and quantify the effect on beliefs, using the assumption that people do not know the true distribution of aggregate economic shocks and learn about it statistically. This is the source of the long-run economic effects. Comparing the resulting outcomes to ones from the same model under the assumption of full knowledge of the distribution (no learning) reveals the extent to which beliefs matter.

But first, we briefly describe the disease spread, the policy reaction and the economic shocks these policies generate.

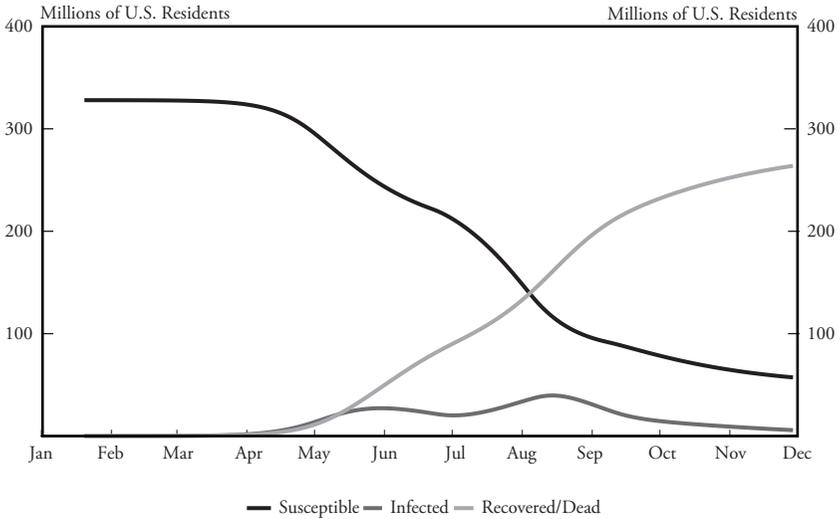
Epidemiology and economic shutdown

Chart 1 illustrates the spread of disease, in both scenarios, as well as the response, which results in capital idling. Recall that scenario 2 has $\zeta_I = 50$, i.e., a policy that is six times less responsive to changes in the infection rate than the $\zeta_I = 300$ policy in scenario 1. As a result, it also has significantly less idle capital and a faster spike in infection rates.

For our purposes, the sufficient statistic in each scenario is the realization for $\tilde{\phi}_{2020}$. In scenario 1, the COVID-19 shock implies $\tilde{\phi}_t = 0.9$, i.e., the loss of value due to obsolescence is equal to 10% of the capital stock. In scenario 2, only 5% of capital is lost to obsolescence: $\tilde{\phi}_t = 0.95$. The target for the initial, transitory impact is line with most forecasts for 2020: a 9% (or 6%) annual decline in GDP.

Chart 1 Disease Spread and Capital Dynamics

Infection Dynamics Scenario 1



Disease R and Capital Scenario 1

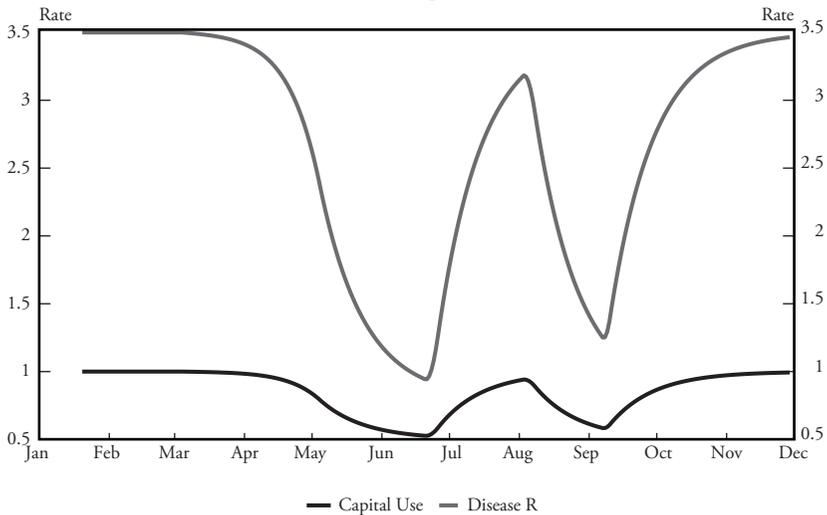
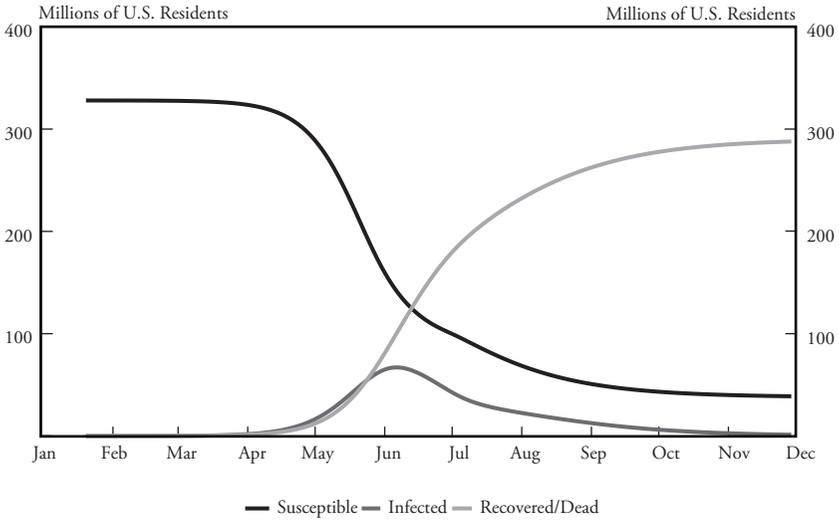
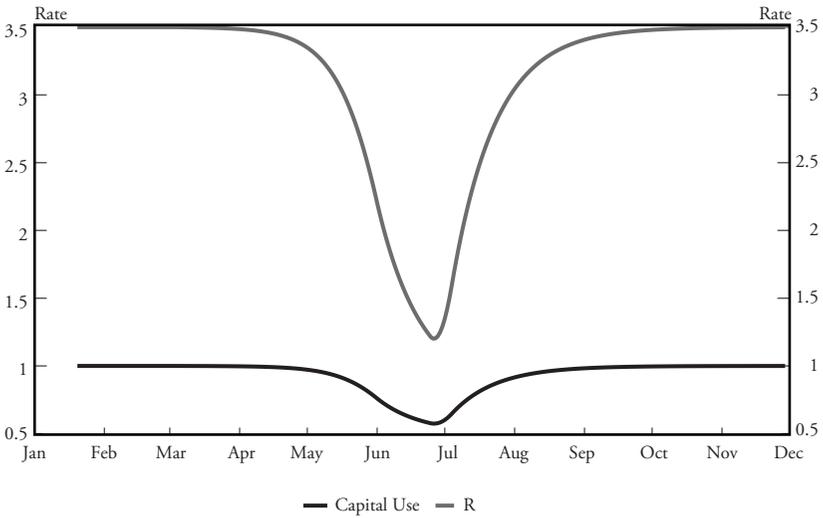


Chart 1 continued

Scenario 2 (less social distance)



Scenario 2



Notes: Parameters listed in Table 1. Scenario 1 uses an aggressive lockdown policy $\zeta_1 = 300$, while scenario 2 uses a more relaxed policy of $\zeta_1 = 50$.

This is likely a conservative estimate for the second quarter of 2020, but more extreme than some forecasts for the entire year.

How much belief scarring?

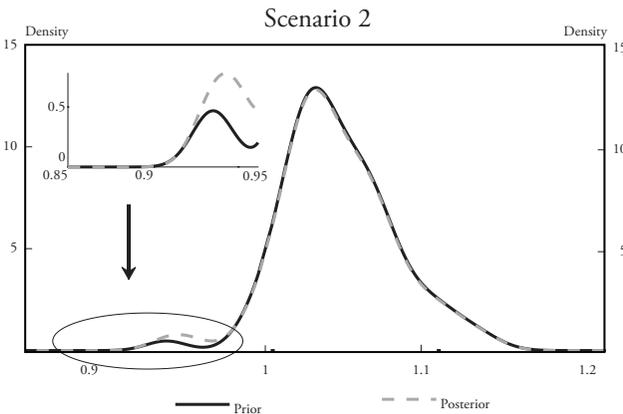
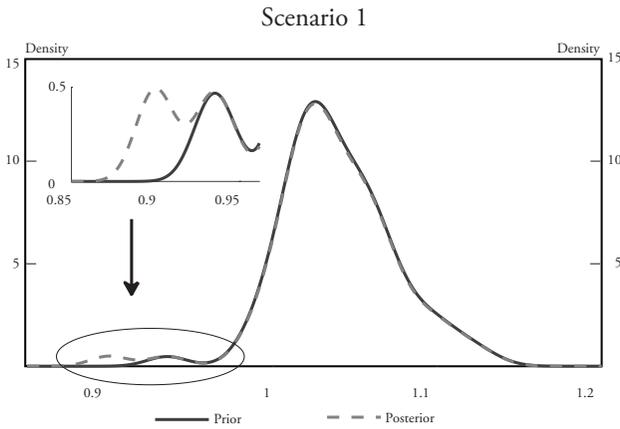
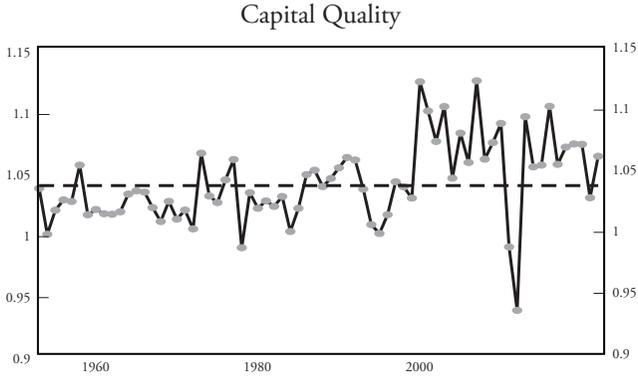
We apply our kernel density estimation procedure to the capital return time series and our two scenarios to construct a sequence of beliefs. In other words, for each t , we construct $\{\hat{g}_t\}$ using the available time series until that point. The resulting estimates for 2019 and 2020 are shown in Chart 2. The differences are subtle. Spotting them requires close inspection where the dotted and solid lines diverge, around 0.90 and 0.95, in scenarios 1 and 2, respectively. They show that the COVID-19 pandemic induces an increase in the perceived likelihood of extreme negative shocks. In scenario 1, the estimated density for 2019 implies near zero (less than 10⁻⁵%) chance of a $\tilde{\phi}_t = 0.90$ shock; the 2020 density attaches a 1-in-70 or 1.4% probability to a similar event recurring.

As the graph shows, for most of the sample period, the shock realizations are in a relatively tight range around 1, but we saw a large adverse realization during the Great Recession of 0.93 in 2009. This reflects the large drops in the market value of non-residential capital stock. The COVID shock is now a second extreme realization of negative capital returns in the last 20 years. It makes such an event appear much more likely.

Effect on GDP

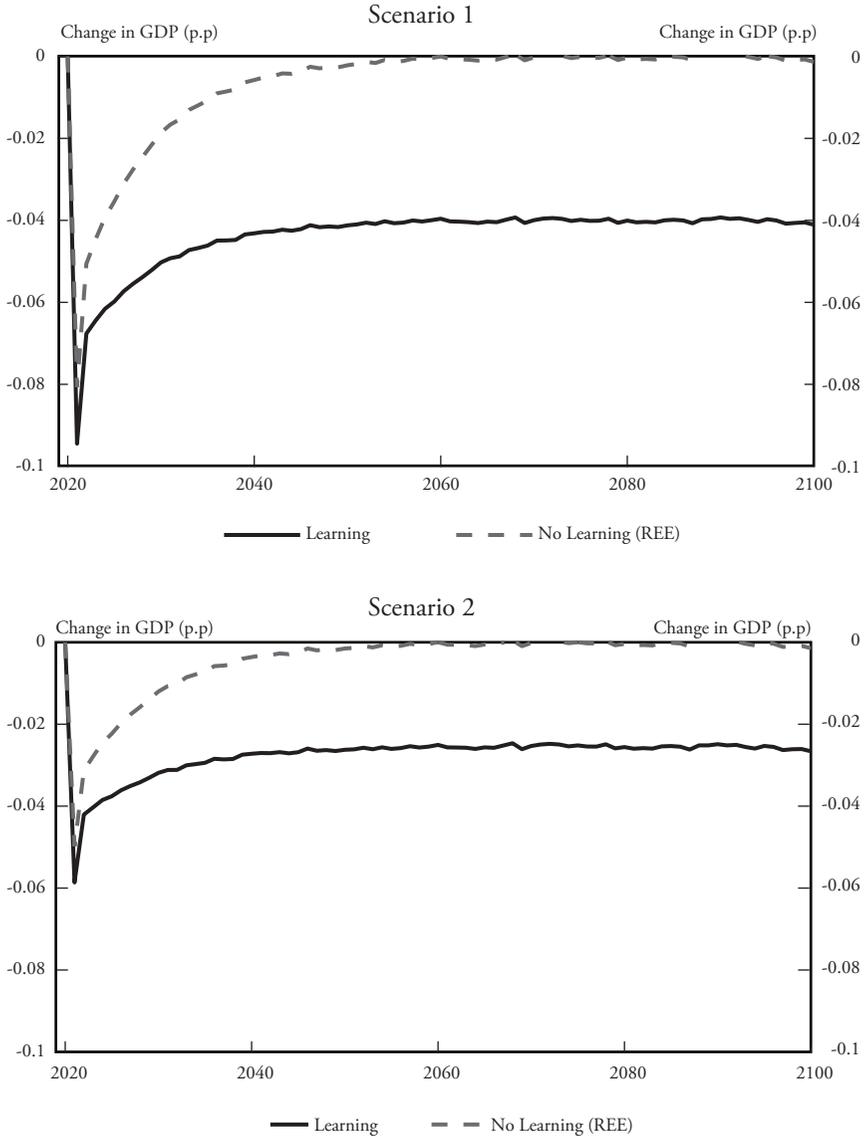
Observing a tail event like the COVID-19 pandemic changes output in a persistent way. Chart 3 compares the predictions of our model for total output (GDP) to an identical model without learning. The units are log changes, relative to the pre-crisis steady-state. In the model without learning, agents are assumed to know the true probability of pandemics. As a result, when they see the COVID crisis, they do not update the distribution. This corresponds to the canonical “rational expectations” assumption in macroeconomics. The model with learning, which uses our real-time kernel density estimation to inform beliefs, generates similar short-term reactions, but worse long-term effects. The post-2020 paths are simulated as follows: each economy is assumed to be at its stochastic steady state in 2019 and is

Chart 2 Beliefs About the Probability Distribution of Outcomes Plotted Before and During the COVID-19 Crisis



Notes: The first panel shows the realizations of ϕ . The second and third panels show the estimated kernel densities for 2019 (solid line) and 2020 (dashed line) for the two scenarios. The subtle changes in the left tail represent the scarring effect of COVID-19.

Chart 3
Output With Scarring of Beliefs (Solid Line) and Without (Dashed Line)



Notes: Units are percentage changes, relative to the pre-crisis steady-state, with 0 being equal to steady state and -0.1 meaning 10% below steady state. Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$, Scenario 2: $\tilde{\phi}_{2020} = 0.95$

subjected to the same 2020 $\tilde{\phi}$ shock; subsequently, sequences of shocks drawn from the estimated 2020 distribution.

The scenarios under learning correspond to what one might call a V-shaped or tilted-V recession: the recovery after the shock has passed is significant but not complete. Note that the drop in GDP on impact is a calibration target—what we are interested in is its persistence, which arguably matters more for welfare. The graph shows that, in Scenario 1, learning induces a long-run drop in GDP of about 4%. The right panel shows a similar pattern but the magnitudes are smaller. Of course, agents also learn from smaller capital obsolescence shocks. These also scar their beliefs going forward. But the scarring is much less, producing only a 3% loss in long-run annual output.

Higher tail risk (i.e., greater likelihood of obsolescence going forwards) increases the risk premium required on capital investments, leading to lower capital accumulation. It is important that these shocks make capital obsolete, rather than just reduce productivity, because obsolescence has a much bigger effect on capital returns than lower productivity does. Labor also contracts, but that is a reaction to the loss of available capital that can be paired with labor. When a chunk of capital becomes maladapted and worthless, that is an order of magnitude more costly to the investor than the temporary decline in capital productivity. Since most of the economic effect works through capital risk deterring investment, that lower return is important to get the economic magnitudes right.

Turning off belief updating

When agents do not learn, both scenarios exhibit quick and complete recoveries, even with a large initial impact. Without the scarring of beliefs, facilities are re-fitted, workers find new jobs, and while the transition is painful, the economy returns to its pre-crisis trajectory relatively quickly. In other words, without belief revisions, the negative shock leads to an investment boom, as the economy replenishes the lost effective capital. While the curvature in utility moderates the speed of this transition to an extent, the overall pattern of a steady recovery back to the original steady state is clear. This is in sharp contrast to

the version with learning. Note that since the no-learning economy is endowed with the same end-of-sample beliefs as the learning model, they both ultimately converge to the same *levels*. But, they start at different steady states (normalized to 0 for each series). This shows that learning is what generates long-lived reductions in economic activity.

Decomposing long-run losses

Next, we perform a simple calculation to put the size of the long-run loss in perspective. Specifically, we use the stochastic discount factor implied by the model to calculate the expected discounted value of the reduction in GDP. These estimates, reported in Table 2, imply that the representative agent in this economy values the cumulative losses between 57% and 90% of the pre-COVID GDP. Most of this comes from the belief scarring mechanism.

Note that the 1-year loss during the pandemic is 6-9% of GDP. The cost of belief scarring is five to six times as large, in both cases. The cost of obsolete capital is about four times as large as the damage done during the pandemic. Chart 4 illustrates the losses each year from the capital obsolescence and belief changes. The area of each of these regions, discounted as one moves to the right in time, is the NPV calculation in Table 2. The one-year cost is a tiny fraction of this total area.

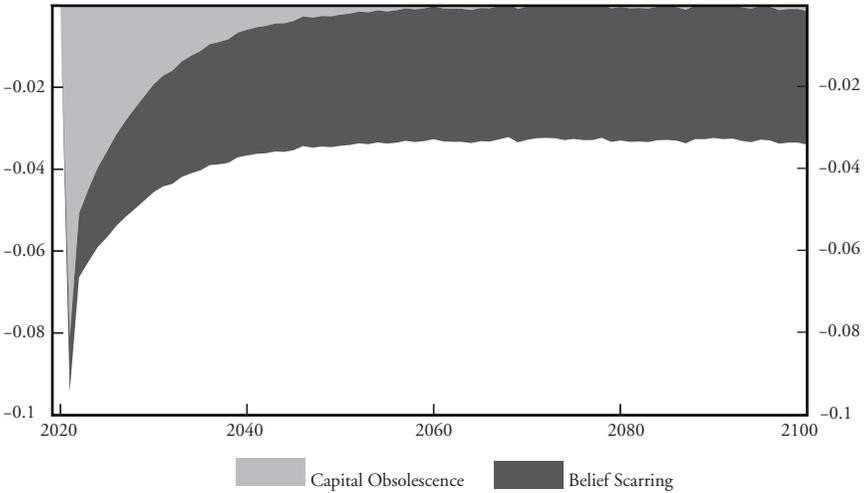
Of course, that calculation misses an important aspect of what we've learned—that pandemics will recur. Since our agents have 70 years of data, during which they've seen one pandemic, they assess the future risk of pandemics to be 1-in-70 initially. That probability declines slowly as time goes on and other pandemics are not observed. But there is also the risk there will be more pandemics, like this. This is not really a result of this pandemic. But that risk of future pandemics is what we should consider if we think about the benefits of public health investments. The pandemic cost going forward, in a world where a pandemic has a 1/70th probability of occurring each year, is given in Chart 5. Note that the risk of future pandemics costs the economy about 7-12% of GDP. This is similar to the one-year cost during the COVID crisis.

Table 2
New Present Value Costs in Percentages of 2019 GDP

Scenario	2020 GDP drop	NPV (Belief Scarring)	NPV (Obsolete capital)
I. Tough	-9%	-52%	-38%
II. Light	-6%	-33%	-24%

Notes: Data comes from FRED. Pre-COVID (post-COVID) data are for Jan. 1, 2020 (July 16, 2020).

Chart 4
Long-Term Costs of the Pandemic
 Scenario 1



Scenario 2

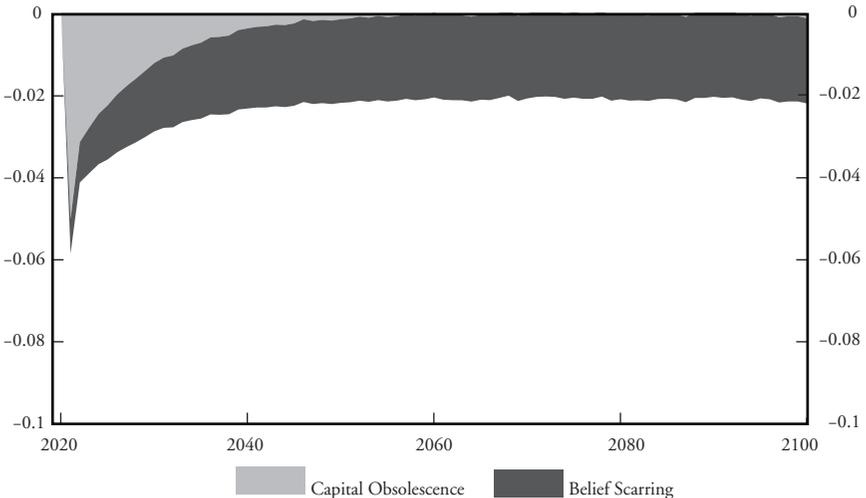
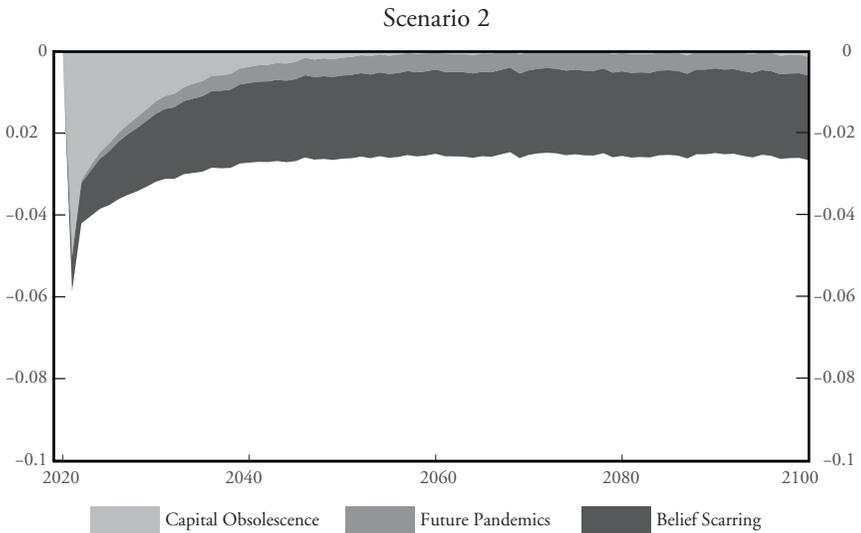
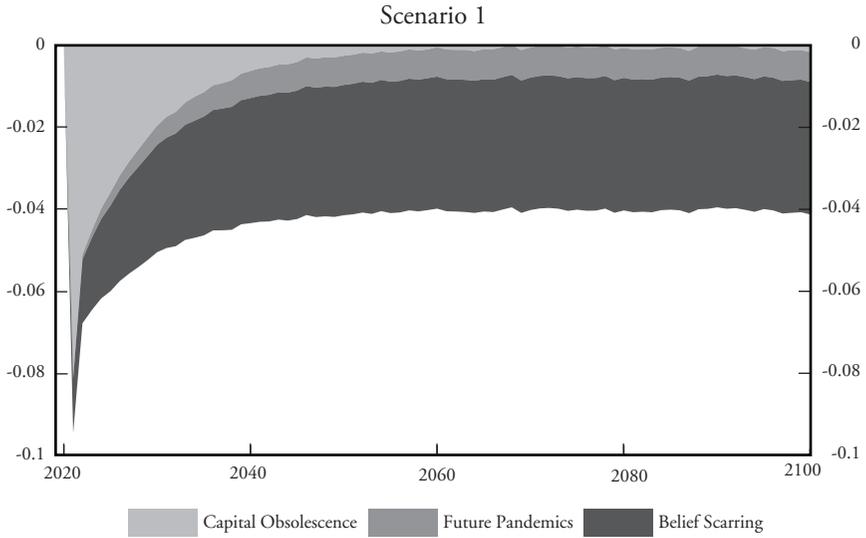


Chart 5 Long-Term Costs of Future Pandemics



Investment and labor

Chart 6 shows the effect of belief changes on investment. When agents do not learn, investment surges immediately (as the economy replenishes the obsolete capital). With learning, investment shows a much smaller surge (starting in 2021), but eventually falls below the pre-COVID levels.

In Chart 7, we see that the initial reaction of labor is milder than for investment, but the bigger differences arise from 2021 onward. When the transitory shock passes, investment surges, to higher than its initial level, to compensate for the obsolescence shock. But labor remains below the pre-COVID levels, reflecting the effect of the scarring effect on the stock of capital and through that on the demand for labor.

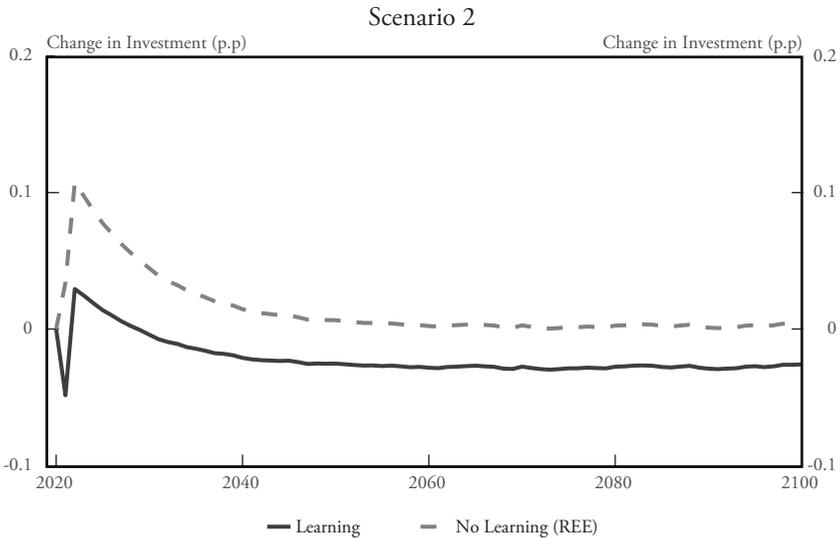
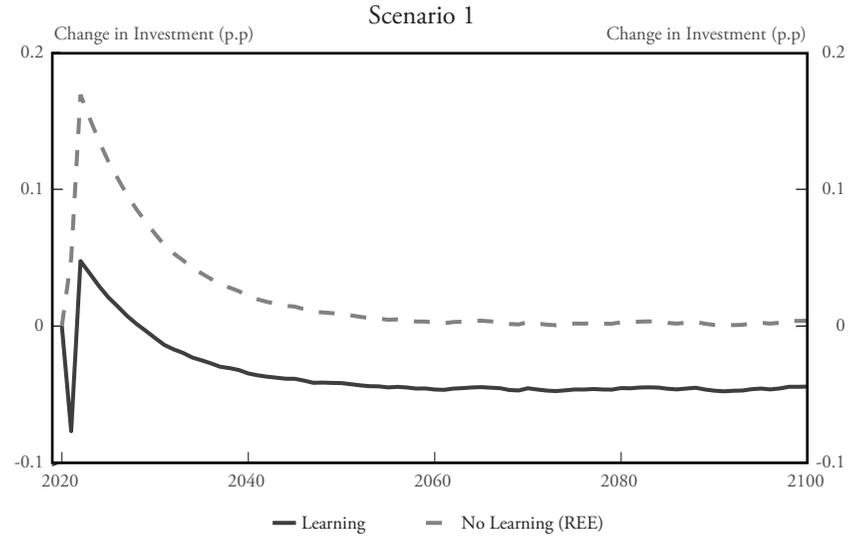
Defaults, riskless rates and credit spreads

The scenarios differ in their short-term implications for default as well. Default spikes only in 2020, the period of impact, returning to average from 2021 onwards. But, the higher default rate in scenario 1 (6% relative to 4% in scenario 2) contributes to greater scarring (since default amplifies obsolescence). This result suggests a role for policy: preventing default/bankruptcy can lead to long-lasting benefits. In Section VI, we present a quantitative analysis of such a policy.

Nearly immediately, after the pandemic passes, default rates in both scenarios return to their original level. While defaults leave permanent scars on beliefs, the defaults themselves are not permanently higher. It is the memory of a transitory event that is persistent.

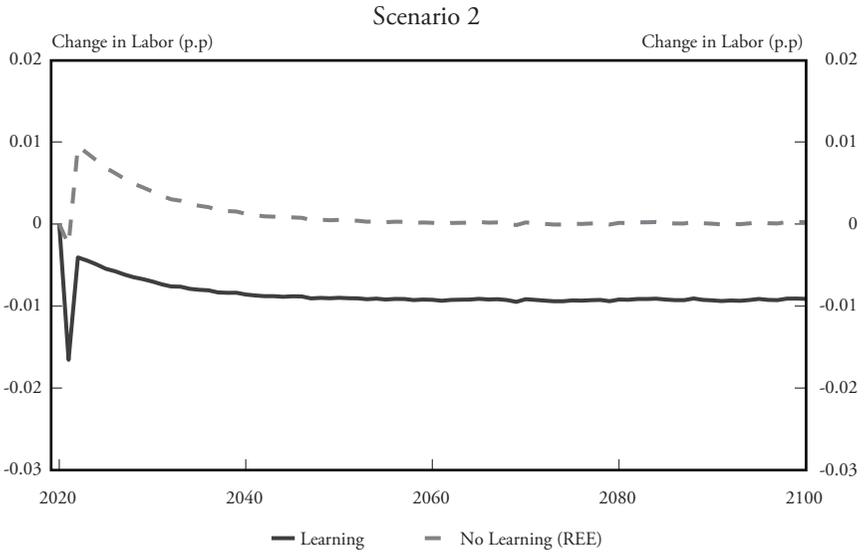
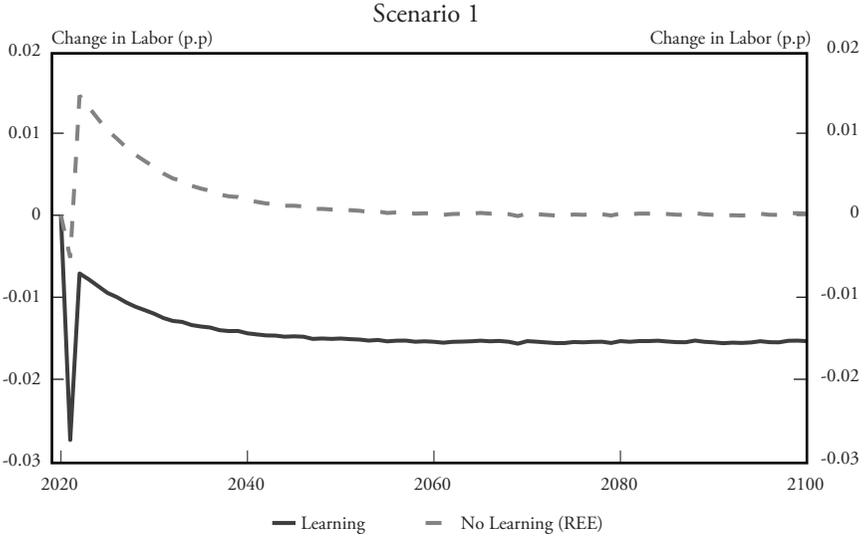
Because defaults were elevated, the pandemic had a large, immediate impact on credit spreads. However, these high spreads were quickly reversed. Some authors have argued that heightened tail risk should inflate risk premia, as well as credit spreads (Hall 2016). While the argument is intuitive, it ignores any endogenous response of discounting, investment or borrowing. A surge in risk triggers disinvestment and deleveraging. Because firms borrow less, this lowers default rates back down, which offsets the increase in the credit spread. We can see this channel at work in the drop in debt and the

Chart 6 Without Belief Scarring, Investment Surges



Notes: Results show average aggregate investment, with scarring of beliefs (solid line) and without (dashed line). Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$, Scenario 2: $\tilde{\phi}_{2020} = 0.95$

Chart 7
Labor With Scarring of Beliefs (Solid Line) and
Without (Dashed Line)



Notes: Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$, Scenario 2: $\tilde{\phi}_{2020} = 0.95$.

Table 3
Changes in Financial Market Variables: Baseline,
Scenarios 1 and 2

	2019	Scenario 1		Scenario 2	
	Baseline	level	change	level	change
Credit Spreads	0.837%	0.842%	+0.5 bps	0.838%	+0.1 bps
Debt	2.75	2.56	-0.19	2.63	-0.12
Default	2.0%	2.0%	0	2.0%	0
Risk free rate	3.66%	3.46%	-20 bps	3.54%	-11 bps
Equity market value	0.44	0.45	+0.01	0.44	0
SKEW	102.7	111.3	+8.6	104.2	+1.5
Third moment $E(R^e - \bar{R}^e)^3$	-1.5	-9.8	-8.3	-2.6	-1.1

Notes: Baseline is the steady state pre-pandemic, under 2019 beliefs. Columns labelled “change” are the raw difference between the long-run average values under 2019 and 2020 beliefs in each scenario. They do not capture any changes that take place along the transition path or during the pandemic. The aggregate market capitalization in our model is the value of the dividend claim times the aggregate capital stock. Third moment is $E(R^e - \bar{R}^e)^3 \times 10^4$, where R^e is the return on equity. The expectation is taken under the risk-neutral measure. For the no-learning model, all changes are zero.

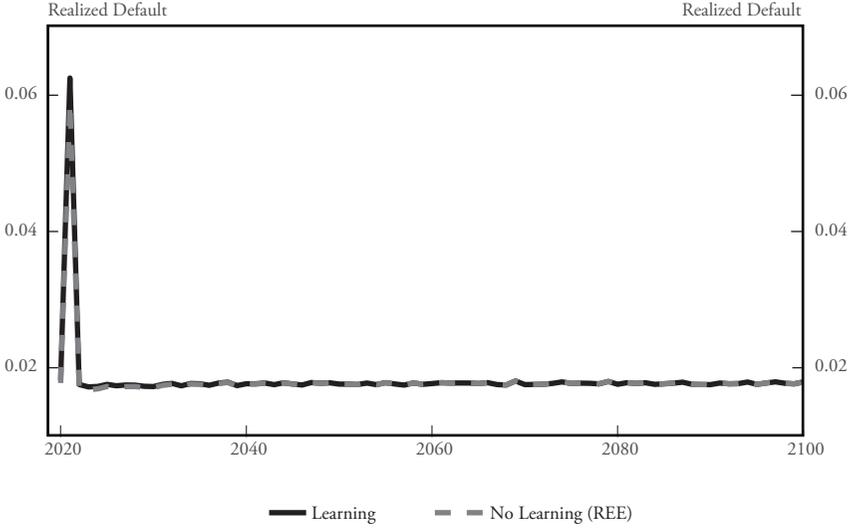
lack of change in long-run defaults (Table 3). The credit spread is the implied interest on risky debt, $1/q_t$ less the risk-free rate r^f . The credit spread in the stochastic steady state under the 2019 belief is less than a basis point higher in the post-pandemic long run. Thus, belief revisions can have significant and long-lived real effects, even if the long-run change in credit spreads is very small.

Equity markets and implied skewness

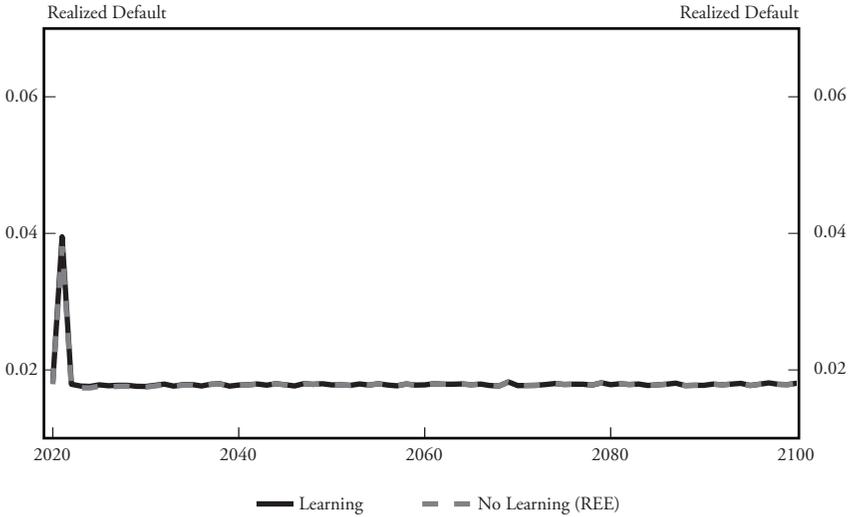
One might think the recent recovery in equity prices appears inconsistent with a persistent rise in tail risk. The model teaches us why this logic is incomplete. When firms face higher tail risk, they also reduce debt, which pushes in the opposite direction as the rise in risk. Furthermore, when firms reduce investment and capital stocks decline, the marginal value of capital rises. Finally, when interest rates and thus future discount rates decline, future equity payments are worth more in present value terms. These competing effects cancel each other out. In our model, the market value of a dividend claim associated with a unit of capital is nearly identical under the post-COVID beliefs than under the pre-COVID ones. In other words, the combined effect of the changes in tail risk and debt reduction is actually mildly positive. While

Chart 8 Realized Default does not Respond Much to Beliefs

Scenario 1



Scenario 2



Notes: Results show with scarring of beliefs (solid line) and without (dashed line) often with the two lines on top of each other. Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$, Scenario 2: $\tilde{\phi}_{2020} = 0.95$.

the magnitudes are not directly interpretable, our point is simply that rising equity valuations are not evidence against tail risk.

If credit spreads and equity premia are not clear indicators of tail risk, what is? For that, we need to turn option prices, in particular out-of-the-money put options on the S&P 500, which can be used to isolate changes in perceived tail risk. A natural metric is the third moment of the distribution of equity returns. The last row of Table 3 reports this object (computed under the risk neutral measure). It shows that the perceived distribution after the shock is more negatively skewed.¹⁶

This might sound inconsistent with the behavior of the SKEW index reported by the CBOE. This showed a short-lived spike at the onset of the pandemic, but recovered quite rapidly. To understand this pattern, note that the SKEW indexes the *standardized* third moment implied by options prices, which is obtained by dividing the third moment from the previous paragraph by the implied standard deviation (or VIX). Tail events typically lead to a spike in market volatility, both realized and implied. This increase in VIX tends to mechanically lower the skewness. More generally, the SKEW index confounds changes in the third moment with the changes in the second moment, which often reflects many other factors. This is the main reason why we focus on the (non-standardized) third moment. As we saw, this measure clearly reveals the persistent change in beliefs and is consistent with evidence from newspapers and surveys in Barrero and Bloom (2020).

V. Liquidity and Interest Rates

In this section, we augment the baseline model to include a liquidity friction. This is motivated by evidence showing liquidity becoming more scarce following the onset of the pandemic—see Boyarchenko et al. (2020). As we will show, a liquidity motive amplifies the effects of tail risk on rates of return for liquid assets, such as Treasuries. This helps bring this dimension of the model's predictions closer to the observed drops in recent months. We also present evidence from bond markets consistent with the rise in liquidity premia.

Recall that, in the baseline model, riskless rates fall in response to higher demand for safe assets. Just as firms react to the increased tail risk by deleveraging, investors would like to protect themselves against low-return states by holding more riskless assets. They cannot all hold more. Therefore, the price increases (the rate of return falls) to clear the market. Table 3 reports the riskless rate falls by 20 bps (10 bps) in scenario 1 (scenario 2). The sign of this change is consistent with what we saw following the onset of the pandemic, but the magnitude is not: interest rates, especially in the Treasury market, fell much more dramatically.

We introduce liquidity considerations using a stylized yet tractable specification, in the spirit of Lagos and Wright (2005).¹⁷ A positive NPV investment opportunity requires liquid funds. Both capital and government bonds provide liquidity (the former only partially). An adverse capital obsolescence shock reduces the value of capital and thus the amount of liquidity it provides. Thus, an increase in the risk of such a shock makes capital liquidity uncertain and raises the value of riskless bonds, which always retain their full, liquid value. Thus, higher tail risk also raises liquidity risk and makes riskless bonds, which serve as liquidity insurance, even more attractive. This channel amplifies the effect on their return and turns out to be quantitatively very large. The increased tail risk brought on by the pandemic, combined with liquidity risk, will turn out to depress interest rates three and a half times as much as in the model without liquidity risk.

Formally, firms are assumed to have access to a profitable intra-period opportunity, yielding a net return of $H(x_t) - x_t$ where x_t is the amount invested. The net return is maximized at $x_t^* = x^*$. But, the firm faces a liquidity constraint: x_t cannot exceed the amount of pledgeable collateral. Formally,

$$x_t \leq a_t + \bar{d}k_t \quad (22)$$

where the parameter \bar{d} indexes the pledgeability of capital and α_t denotes a riskless, fully liquid asset. This can be interpreted narrowly as government bonds,¹⁸ but it could also be thought of as the total liquidity available from other sources. Note that the liquidity value of capital

is a function of *effective* capital, i.e., net of obsolescence. As a result, shocks to capital obsolescence influence the availability of liquidity.

The supply of the liquid asset is assumed to be fixed at \bar{a} . Thus, the amount invested in the opportunity in t is given by $x_t = \min(x^*, \bar{a} + \bar{d}k_t)$. The liquidity premium is the marginal value (in units of consumption) of an additional unit of pledgeable collateral:

$$\mu_t = H'(x_t) - 1. \quad (23)$$

The return on government bonds, i.e., the liquid asset, is characterized

$$\frac{1}{R^g} = \mathbb{E}_t [M_{t+1}(1 + \mu_{t+1})]. \quad (24)$$

The final model alternation is that the liquidity premium shows up in the first term of the optimality condition for capital (16), which becomes $\mathbb{E}_t [M_{t+1}(R_{t+1}^k + \bar{d}\mu_{t+1})]$.

Parameterization

To set values for the liquidity parameters, we follow the strategy in Kozłowski et al. (2019). We use the following functional form for the benefit to invest on liquid assets: $H(x) = 2\iota\sqrt{x} - \xi$. The parameter that governs how much of capital is a pledgeable, liquid asset, \bar{d} , is set to 0.16 to match the ratio of short-term obligations of U.S. non-financial corporations to the capital stock in the Flow of Funds. The liquid asset supply $\bar{a} = 0.8$ and the return parameter $\iota = 1.4$ are chosen so that the ratio of liquid assets to capital is 0.08 and their return in the pre-COVID steady state equals 2%. Finally, the parameter $\xi = 1.94$ is set so the net return of the project is close to zero (on average) in the pre-COVID steady state.

Riskless rates with liquidity premia

The purpose of this extension was to explore how liquidity considerations affect scarring-induced changes in riskless rates. To evaluate this, we compute riskless rates in the stochastic steady states associated with the pre- and post-COVID beliefs. These are presented in Table 4. The model predicts that the yield on liquid bonds drops by 67 bps in the new steady state. In contrast, the return on a riskless but

Table 4
Implications for Interest Rates with Liquidity Frictions,
Model vs. Data

	Pre-COVID	Post-COVID	Chg
Model			
Riskless rate (liquid) $R_f^l - 1$	2.12%	1.46 %	-67 bps
Riskless rate (illiquid)	4.97%	4.89 %	-8 bps
Data			
1y Treasury yield (nominal)	1.56%	0.14%	-142 bps
5y Treasury yield (nominal)	1.67%	0.28%	-139 bps
5y forward rate (real)	-0.09%	-0.98%	-89 bps
AAA Yield	2.53%	1.48%	-105 bps
AAA Spread (rel. to 5y Treasury)	0.86%	1.20%	34 bps

Notes: Data comes from FRED. Pre-COVID (post-COVID) data are for Jan. 1, 2020 (July 16, 2020).

completely illiquid asset falls only by 8 bps: in other words, the liquidity premium rises by 59 bps.

The table also shows changes in various market interest rates between January and July 2020. The yield on the one-year and five-year Treasuries were almost 1.4% lower in July 2020 (relative to the beginning of the year). Note that these are not directly comparable with the model numbers. The latter compare steady-states and so are most appropriately thought of as long-run predictions while the current data obviously reflect short-term, more transitory considerations as well. We therefore construct a proxy for the long-run rates using forward rates implied by the Treasury yield and long-term inflation expectations. Specifically, we use the instantaneous rate five years forward from the Treasury yield curve and 5y5y inflation expectations¹⁹ to calculate the change in long-term real rates. This shows a decline of about 89 bps, smaller than short-term rates and closer to the model's predictions.

Next, the table also reports the change in the yield on AAA corporate bonds. These securities carry very little default risk, but are not as liquid as Treasuries. As a result, the yield spread on these bonds relative to Treasuries is often viewed as a proxy for liquidity premia—see,

e.g., Krishnamurthy and Vissing-Jorgensen (2012) and Del Negro et al. (2017). In recent months, this spread rose by 34 bps, consistent with increased liquidity scarcity. The model liquidity premium reported in the table shows a larger rise. This is to be expected since the model object is defined as the spread of a *completely* illiquid security whereas AAA bonds are probably partially liquid.

Chart 9 shows the time path of the natural rate of interest. Notice that the short-run fluctuations are much larger than the long-term effects reported in the table. This is consistent with short-term market disruptions that are now settling down.

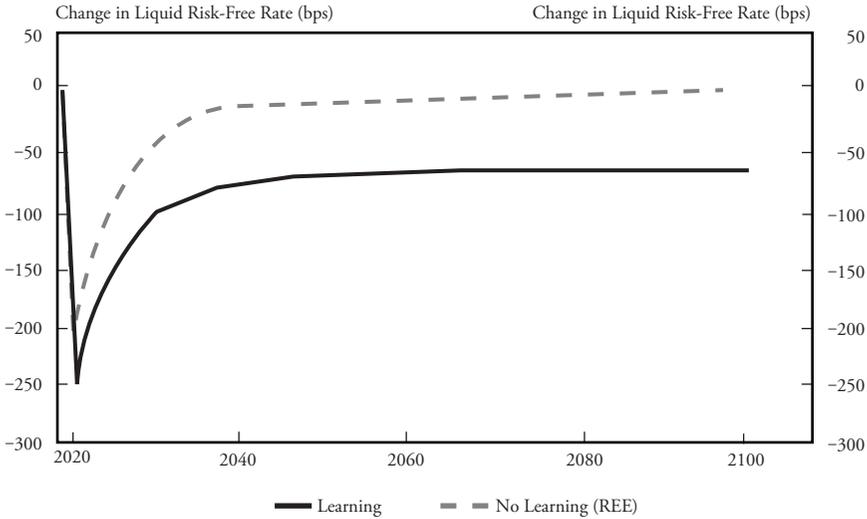
Finally, in interpreting recent data, it is worth pointing out that the last few months have seen unprecedented policy interventions in bond markets, which almost certainly have contributed to the drop in interest rates on both liquid and illiquid assets—see Boyarchenko et al. (2020). Our analysis completely abstracts from such interventions²⁰ so it is perhaps not too surprising that the model underpredicts the fall in interest rates. Overall, these results suggest a quantitatively meaningful role for the belief scarring mechanism in the recent behavior of interest rates.

VI. The Role of Financial Policy

The COVID-19 pandemic has sparked an unprecedented policy response. These responses fall into three broad categories: social distancing and other mobility restrictions, transfers to households and financial assistance to firms. We explored the consequences of more lax social distance policy in constructing scenarios for our baseline results. Transfers to households has an important role to mitigate the economic fallout, but does not directly affect productive capacity, the key object in our analysis. Financial assistance to firms, on the other hand, can help the economy maintain productive capacity, for example by preventing widespread bankruptcies. In our setting, such a policy would have beneficial long-run effects as well, since they mitigate the consequences of belief scarring. In this section, we use our baseline model to quantify these long-run benefits. We find that the longer-term effects of a policy of debt relief are as much as 10 times larger than the short-run effects.

Chart 9

Belief Scarring Lowers Riskless Rate in Long Run



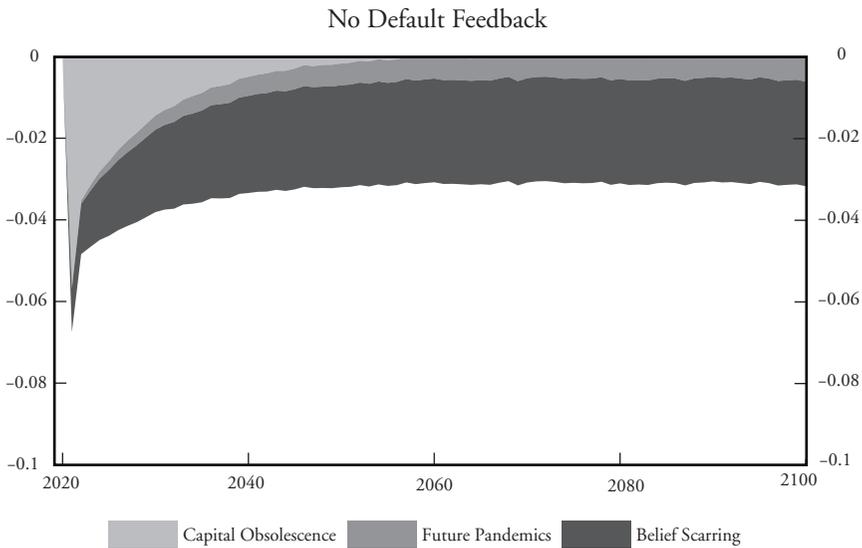
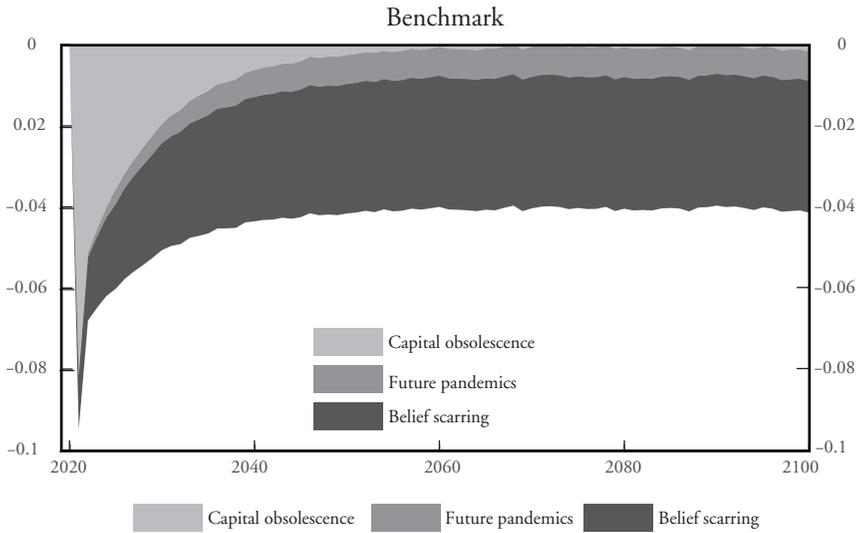
Notes: Results show the return on a riskless asset, in scenario 1, with scarring of beliefs (solid line) and without (dashed line). Common parameters listed in Table 1. Scenario-specific parameters are Scenario 1: $\tilde{\phi}_{2020} = 0.90$.

The need for policy intervention in the model stems from the presence of debt and the associated risk of bankruptcy. Bankruptcy is socially costly because it exacerbates capital obsolescence. Therefore, we model financial policy as designed to prevent/limit bankruptcies by reducing firms' effective debt. This could take the form of the government or other policymaker buying up the debt from private creditors or offering direct assistance to firms. Before examining the effects of such a policy, we perform a simple exercise to quantify the costs of bankruptcy in our baseline model. This is the cost that financial policy might plausibly remedy.

Effect of the default-obsolescence feedback

To understand the role of this feedback rule, suppose that obsolescence is entirely exogenous, i.e., it does not vary with default. This amounts to setting $\mu = 0$ in (11). Chart 10 shows the GDP impact of the COVID-19 shock under our benchmark specification (in the left panel) and without default feedback (in the right panel). The broad patterns are similar with belief revisions accounting for a significant portion of the impact, but the magnitudes are slightly smaller in the

Chart 10
Default Feedback Increases Long-Run Effects



Note: Results show with scarring of beliefs (solid line) and without (dashed line) often with the two lines on top of each other. Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$
 Scenario 2: $\tilde{\phi}_{2020} = 0.95$.

Table 5
Firm Financial Assistance Policy: No Assistance vs.
10% Debt Reduction

	No assistance	10% debt reduction	Benefit
GDP drop in 2020	-9%	-8%	1%
NPV of long-term output loss			
from belief scarring	-52%	-45%	7%
from obsolescence	-38%	-34%	4%

Notes: Results are for scenario 1 ($\tilde{\phi}_{2020} = 0.90$). Numbers shown are in percentages of the pre-COVID steady-state GDP.

right panel (GDP falls by just under 4% in the long run, relative to a 5% drop in the benchmark). This difference between the two panels is the effect of the default-obsolescence feedback.

VI.i. Financial Assistance Policy

We consider a simple policy that prevents bankruptcy by reducing firms debt burden, specifically a reduction in each firm's debt by 10%. This in turn mitigates the effective capital obsolescence and consequently beliefs are slightly less pessimistic going forward. We then simulate the model with these new beliefs and calculate the short- and long-term GDP effects, reported in Table 5.

The table shows that financial assistance of this magnitude only saves 1% of GDP in 2020. From that metric alone, one might judge the cost of the policy to be too high.²¹ However, preventing bankruptcies in the short run also helps reduce losses over time. The present discounted value of those savings are worth 11% of 2019 GDP. Of that 11%, 7% comes from ameliorating belief scarring and another 4% comes from the direct effects of limiting capital obsolescence. This exercise shows that considering the long-run consequences can significantly change the cost-benefit analysis for financial policies aimed at assisting firms.

VII. What If We Had Seen a Pandemic Like This Before?

In our benchmark analysis, pre-COVID beliefs were formed using data that did not witness a pandemic (though it did have other tail

events like the 2007-09 Great Recession). But, pandemics have occurred before—Jorda et al. (2020) identify 12 pandemics (with greater than 100,000 deaths) going to the 14th century. This raises the possibility that economic agents in 2019 had some awareness of these past tail events and believed that they could happen again. To understand how this might change our results, we assume that the pre-COVID data sample includes the 1918 episode. Unfortunately, we do not have good data on capital utilization and obsolescence during that period,²² so we simply use the time series for the capital return shock $\tilde{\phi}_t$ from 1950-2020 as a proxy for the $\tilde{\phi}_t$ series from 1880-1949. In other words, we are asking: What if we had seen all of this unfold exactly the same way before?

The previous data does not change the short-term impact of the shock. But, it does cut the long-term effect of it in half. Just before the pandemic of 2020 struck, our data tells us that there has been one pandemic in nearly 140 years. We assess the probability to be about 1-in-140. After 2020, we saw two pandemics in 141 years. Therefore, we revise our perceived probability from 1-in-140 to 2-in-141. That is about half the change in probability, relative to the original model where the probability rose from zero to 1-in-70.

But considering data from so long ago does raise the question of whether it is perceived as less relevant. There is a sense that the world has changed, institutions are stronger, science has advanced, in ways that alter the probability of such events. Such gradual change might logically lead one to discount old data.

In a second exercise, we assume that agents discount old data at the rate of 1% per year. In this case, two forces compete. The presence of the 1918 events in the sample reduces the surprise of the new pandemic as before, albeit with a much smaller weight. The countervailing force is that when old data is down-weighted, new data is given a larger weight in beliefs. The larger role of the recent pandemic in beliefs going forward makes belief scarring stronger for the next few decades. These forces more or less cancel each other out leaving the net results indistinguishable, in every respect, from the original results with data only from 1950.

Of course, more recently, we saw SARS, MERS and Ebola arise outside the United States. Other countries may have learned from these episodes. But the lack of preparation and slow response to events unfolding in China suggests that U.S. residents and policymakers seem to have inferred only that diseases originating abroad stay outside the U.S. borders.

VIII. Conclusion

No one knows the true distribution of shocks to the economy. Macroeconomists typically assume that agents in their models know this distribution, as a way to discipline beliefs. For many applications, assuming full knowledge has little effect on outcomes and offers tractability. But for unusually large events, like the current crisis, the difference between knowing these probabilities and estimating them with real-time data can be large. We argue that a more plausible assumption for these phenomena is to assume that agents do the same kind of real-time estimation along the lines of what an econometrician would do. This introduces new, persistent dynamics into a model with otherwise transitory shocks. The essence of the persistence mechanism is this: once observed, a shock (a piece of data) stays in one's data set forever and therefore persistently affects belief formation. The less frequently similar data is observed, the larger and more persistent the belief revision.

When we quantify this mechanism, our model's predictions tell us that the ongoing crisis will have large, persistent adverse effects on the U.S. economy, far greater than the immediate consequences. Preventing bankruptcies or permanent separation of labor and capital, could have enormous consequences for the value generated by the U.S. economy for decades to come

Authors' Notes: The views expressed are those of the authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors. We thank Kenneth Rogoff for a generous and insightful discussion, as well as Dean Corbae and Pablo D'Erasmus for sharing data on corporate defaults and Marco Del Negro and Andrea Tambalotti for data.

Appendix A Solution

A.A. Equilibrium Characterization

An equilibrium is the solution to the following system of equations:

$$1 = \mathbb{E}M_{t+1}[R_{t+1}^k]J^k(\underline{v}_t) \tag{25}$$

$$R_{t+1}^k = \frac{(1-\alpha)\phi_{t+1}^{\alpha+\nu}\hat{k}_{t+1}^\alpha I_{t+1}^{1-\alpha} + (1-\delta)\phi_{t+1}\hat{k}_{t+1}}{\hat{k}_{t+1}} \tag{26}$$

$$\frac{1-\gamma}{\gamma} \frac{c_t}{1-I_t} = \frac{(1-\alpha)y_t}{I_t} \tag{27}$$

$$(1-\theta)\mathbb{E}_t[M_{t+1}\underline{v}_t f(\underline{v}_t)] = \left(\frac{\chi-1}{\chi}\right)\mathbb{E}_t[M_{t+1}(1-F(\underline{v}_t))] \tag{28}$$

$$c_t = \phi_t^{\alpha+\nu}\hat{k}_t^\alpha I_t^{1-\alpha} + (1-\delta)\phi_t\hat{k}_t - \hat{k}_{t+1} \tag{29}$$

$$U_t = \left[(1-\beta)(u(c_t, I_t))^{1-\psi} + \beta\mathbb{E}(U_{t+1}^{1-\eta})^{\frac{1-\psi}{1-\eta}} \right]^{\frac{1}{1-\psi}} \tag{30}$$

where

$$\ln \phi_t = \ln \tilde{\phi}_t - \mu F(\underline{v}_t)^{1-\sigma}$$

$$\underline{v}_t = \frac{Ie\nu_{t+1}}{R_{t+1}^k}$$

$$J^k(\underline{v}_t) = 1 + (\chi-1)\underline{v}_t(1-F(\underline{v}_t)) + (\chi\theta-1)h(\underline{v}_t)$$

$$M_{t+1} = \left(\frac{dU_t}{dc_t}\right)^{-1} \frac{dU_t}{dc_{t+1}} = \beta \left[\mathbb{E}(U_{t+1}^{1-\eta}) \right]^{\frac{\eta-\psi}{1-\eta}} U_{t+1}^{\psi-\eta} \left(\frac{u(c_{t+1}, I_{t+1})}{u(c_t, I_t)} \right)^{-\psi}$$

A.B. Solution Algorithm

To solve the system described above at any given date t (i.e., after any observed history of $\tilde{\phi}_t$), we recast it in recursive form with grids for the aggregate state (\hat{k}) and the shocks $\tilde{\phi}$. We then use an iterative procedure:

- Estimate \hat{g} on the available history using the kernel estimator.
- Start with a guess (in polynomial form) for $u(\hat{k}, \tilde{\phi}), c(\hat{k}, \tilde{\phi}), l(\hat{k}, \tilde{\phi})$.
- Solve (25) – (28) for \hat{k}', lev', l using a non-linear solution procedure.
- Verify/update the guess for U, c, l using (29)–(30) and iterate until convergence.

Endnotes

¹See e.g. Summers FT, 2020, Krugman 2020, Reinhart and Rogoff 2020 and Cochrane 2020.

²Other learning papers in this vein include papers on news shocks, such as Beaudry and Portier 2004, Lorenzoni 2009, Veldkamp and Wolfers 2007, uncertainty shocks, such as Jaimovich and Rebelo 2006, Bloom et al. 2018, Nimark 2014 and higher-order belief shocks, such as Angeletos and La'O 2013 or Huo and Takayama 2015.

³As $\kappa_t \rightarrow 0$, the CDF of the kernel converges $\hat{G}_t^0(\tilde{\phi}) = \frac{1}{n_t} \sum_{s=0}^{n_t-1} \mathbf{1}\{\tilde{\phi}_{t-s} \leq \tilde{\phi}\}$. Then, for any $\tilde{\phi}$ and any $j \geq 1$

$$\mathbb{E}_t \left[\hat{G}_{t+j}^0(\tilde{\phi}) \middle| \mathcal{I}_t \right] = \mathbb{E}_t \left[\frac{1}{n_t + j} \sum_{s=0}^{n_t+j-1} \mathbf{1}\{\tilde{\phi}_{t+j-s} \leq \tilde{\phi}\} \middle| \mathcal{I}_t \right] = \frac{n_t}{n_t + j} \hat{G}_t^0(\tilde{\phi}) + \frac{j}{n_t + j} \mathbb{E}_t \left[\mathbf{1}\{\tilde{\phi}_{t+1} \leq \tilde{\phi}\} \middle| \mathcal{I}_t \right]$$

Thus, future beliefs are, in expectation, a weighted average of two terms—the current belief and the distribution from which the new draws are made. Since our best estimate for the latter is the current belief, the two terms are exactly equal, implying $\mathbb{E}_t \left[\hat{G}_{t+j}^0(\tilde{\phi}) \middle| \mathcal{I}_t \right] = \hat{G}_t^0(\tilde{\phi})$.

⁴See Hansen 2015.

⁵This is consistent with the U.S. official policy on re-opening (CDC 2020). Note that individual optimal choice to social distance are also included in this “policy.” These optimal choices look similar. See Kaplan et al. 2020.

⁶The i.i.d. assumption also has empirical support. In the next section, we use macro data to construct a time series for $\tilde{\phi}_t$. We estimate a statistically insignificant autocorrelation of 0.15.

⁷This is a natural assumption: with a continuum of firms and a stationary shock process, firms can learn the complete distribution of any idiosyncratic shocks after one period.

⁸In our baseline specification, default does not destroy resources—the penalty is purely private. This is not crucial—it is straightforward to relax this assumption by assuming that all or part of the cost of the default represents physical destruction of resources.

⁹The subsidy is assumed to be paid by a government that finances it through a lump-sum tax on the representative household.

¹⁰To simulate the COVID-19 pandemic, we run the epidemiology model from Section II.i for one year and use the predicted capital obsolescence as the realized shock for 2020. For more details, see Section III.

¹¹These are series FL102010005 and FL102010115 from Flow of Funds.

¹²Our results are robust to alternative measures of nominal price changes, e.g., computed from the price index for GDP or Personal Consumption Expenditure.

¹³The Moody's data is for rated firms and shows a historical average default rate of 1% (our calibration implies a default rate of 2%), probably reflecting selection. Accordingly, we scaled the Moody's estimates by a factor of 2 while performing this calculation. We also used estimates of exit and bankruptcy rates from Corbae and D'Erasmus 2017 and found broadly similar results.

¹⁴This is in line with the target in Khan et al. 2017, though a bit higher than the one in Gourio 2013. We verified that our quantitative results are not sensitive to this target.

¹⁵Section VI studies a version without default amplification and finds that it generates similar patterns, albeit with slightly smaller magnitudes, as our benchmark economy.

¹⁶It is straightforward to compute this from the SKEW and VIX indices reported by the CBOE. The third central moment under the risk-neutral measure is $E(R^e - \bar{r}^e)^3 = \frac{100 - SKEW_t}{10} \cdot VIX_t^3$. This calculation reveals that, between February and May 2020, the market implied third moment also became significantly more negative (from -0.04 to -0.09).

¹⁷See also Kozłowski et al. 2019.

¹⁸For concreteness, we adopt this assumption in our analysis. The bonds are issued by a government, which balances its budget with lump-sum taxes/transfers.

¹⁹5y5y inflation expectations are the expectations of inflation over the five-year period, starting five years from today. Source: FRED, Federal Reserve Bank of St. Louis. The series tickers are THREEFF5 and T5YIFR, respectively.

²⁰We do evaluate the effects of a financial policy in the following section.

²¹Strictly speaking, in our model with a representative agent and lump-sum taxes, there is no real cost to implementing this policy. But, obviously, in a more realistic setting with heterogeneity and distortionary taxation, taking over 10% of corporate debt would entail substantial costs/dead-weight losses.

²²See Correia et al. 2020 and Velde 2020 for analysis of the economic effects of the 1918 over the short- to medium-term.

References

- Alvarez, F.E., D. Argente and F. Lippi. 2020, April. "A Simple Planning Problem for COVID-19 Lockdown," Working Paper 26981, National Bureau of Economic Research.
- Angeletos, G.-M., and J. La'O. 2013. "Sentiments," *Econometrica*, 81(2), 739-779.
- Atkeson, A. 2020, March. "What Will Be the Economic Impact of COVID-19 in the U.S.? Rough Estimates of Disease Scenarios," Working Paper 26867, National Bureau of Economic Research.
- Baker, S., N. Bloom, S. Davis and K. Kost. 2019, March. "Policy News and Stock Market Volatility," working paper.
- Barrero, J., and N. Bloom. 2020. "Economic Uncertainty and the Recovery," in *Navigating the Decade Ahead: Implications for Monetary Policy*, Federal Reserve Bank of Kansas City, Jackson Hole Economic Policy Symposium, Kansas City, Missouri, Aug. 27-28.
- Beaudry, P., and F. Portier. 2004. "An Exploration into Pigou's Theory of Cycles," *Journal of Monetary Economics*, 51, 1183-1216.
- Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten and S.J. Terry. 2018. "Really Uncertain Business Cycles," *Econometrica*, 86(3), 1031-1065.
- Boyarchenko, N., A. Kovner and O. Shachar. 2020. "It's What You Say and What You Buy: A Holistic Evaluation of the Corporate Credit Facilities," Technical report, Federal Reserve Bank of New York, Staff Report.
- Carvalho, V.M., S. Hansen, Á. Ortiz, J.R. Garcia, T. Rodrigo, S. Rodriguez Mora and P. Ruiz de Aguirre. 2020. "Tracking the Covid-19 Crisis with High-Resolution Transaction Data."
- Cochrane, J. 2020. "What Shape Recovery?" Available at <https://johnhcochrane.blogspot.com/2020/04/what-shape-recovery.html> (2020/04/06).
- Cogley, T., and T. Sargent. 2005. "The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty," *Review of Economic Dynamics*, 8, 528-563.
- Corbae, D., and P. D'Erasmus. 2017, June. "Reorganization or Liquidation: Bankruptcy Choice and Firm Dynamics, Working Paper 23515, National Bureau of Economic Research.
- Correia, S., S. Luck and E. Verner. 2020. "Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu," SSRN Working Paper Series.

- CPPI. 2020. Commercial Property Price Index. Technical report, Green Street Advisors. Accessed: May 15, 2020.
- Del Negro, M., D. Giannone, M. Giannoni and A. Tambalotti. 2017. "Safety, Liquidity, and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, 235-294.
- Eaton, J., and M. Gersovitz. 1981. "Debt with Potential Repudiation: Theoretical and Empirical Analysis," *The Review of Economic Studies*, 289-309.
- Eichenbaum, M.S., S. Rebelo and M. Trabandt. 2020, March. "The Macroeconomics of Epidemics," Working Paper 26882, National Bureau of Economic Research.
- Elenev, V., T. Landvoigt and S. Van Nieuwerburgh. 2020. "Can the COVID Bailouts Save the Economy?" *COVID Economics*, 17.
- Fajgelbaum, P.D., E. Schaal and M. Taschereau-Dumouchel. 2017, 05. "Uncertainty Traps," *The Quarterly Journal of Economics*, 132(4), 1641-1692.
- Farboodi, M., G. Jarosch and R. Shimer. 2020, April. "Internal and External Effects of Social Distancing in a Pandemic," Working Paper 27059, National Bureau of Economic Research.
- Gangestad, S., and D. Buss. 1993. "Pathogen Prevalence and Human Mate Preferences," *Ethology and Sociobiology*, 14(2), 89-96.
- Gourio, F. 2013. "Credit Risk and Disaster Risk," *American Economic Journal: Macroeconomics*, 5(3), 1-34.
- Gourio, F. 2012. "Disaster Risk and Business Cycles," *American Economic Review*, 102(6), 2734-2766.
- Hall, R.E. 2016. "Macroeconomics of Persistent Slumps," in *Handbook of Macroeconomics*, Vol. 2, pp. 2131-2181. Elsevier.
- Hansen, B. E. 2015. *Econometrics*.
- Hansen, L. 2007. "Beliefs, Doubts and Learning: Valuing Macroeconomic Risk," *American Economic Review*, 97(2), 1-30.
- Huo, Z., and N. Takayama. 2015. *Higher Order Beliefs, Confidence, and Business Cycles*.
- Jaimovich, N., and S. Rebelo. 2006. "Can News About the Future Drive the Business Cycle?" *American Economic Review*, 99(4), 1097-1118.
- Johannes, M., L. Lochstoer and Y. Mou. 2016. "Learning About Consumption Dynamics," *The Journal of Finance*, 71(2), 551-600.

- Jorda, O., S. Singh and A. Taylor. 2020. "Longer-Run Economic Consequences of Pandemics," *COVID Economics*, 1, 1-15.
- Kaplan, G., B. Moll and G. Violante. 2020. "Pandemics according to HANK."
- Kermack, W.O., and A.G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics," *Proceedings of the Royal Society of London*, Series A, 115(772), 700-721.
- Khan, A., T. Senga and J.K. Thomas. 2017. "Default Risk and Aggregate Fluctuations in an Economy with Production Heterogeneity," 2017 Meeting Papers 889, Society for Economic Dynamics.
- Koren, M., and R. Pető. 2020. "Business Disruptions from Social Distancing," *COVID Economics*, 2, 13-31.
- Kozeniauskas, N., A. Orlik and L. Veldkamp. 2014. "What Are Uncertainty Shocks?" *Journal of Monetary Economics*, 100, 1-15.
- Kozlowski, J., L. Veldkamp and V. Venkateswaran. 2020. "The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation," *Journal of Political Economy*, 128(8), 2839-2879.
- Kozlowski, J., L. Veldkamp and V. Venkateswaran. 2019. "The Tail that Keeps the Riskless Rate Low," *National Bureau of Economic Research, Macroeconomics Annual*, 33(1), 253-283.
- Krishnamurthy, A., and W. Li. 2020, May. "Dissecting Mechanisms of Financial Crises: Intermediation and Sentiment," Working Paper 27088, National Bureau of Economic Research.
- Krishnamurthy, A., and A. Vissing-Jorgensen. 2012. "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120(2), 233-267.
- Krugman, P. 2020. "Notes on the Coronacoma," available at <https://www.nytimes.com/2020/04/01/opinion/notes-on-the-coronacoma-wonkish.html> (April 1, 2020).
- Lagos, R., and R. Wright. 2005. "A Unified Framework for Monetary Theory and Policy Analysis," *Journal of Political Economy*, 113(3), 463-484.
- Leibovici, F., A.M. Santacreu and M. Famiglietti. 2020. "How the Impact of Social Distancing Ripples through the Economy," available at <https://www.stlouisfed.org/on-the-economy/2020/april/impact-social-distancing-ripples-economy> (April 7, 2020).
- Loehle, C. 1995. "Social Barriers to Pathogen Transmission in Wild Animal Populations," *Ecology*, 76(2), 326-335.
- Lorenzoni, G. 2009. "A Theory of Demand Shocks," *American Economic Review*, 99(5), 2050-2084.

- Ludvigson, S.C., S. Ma and S. Ng. 2020, April. "Covid19 and the Macroeconomic Effects of Costly Disasters," Working Paper 26987, National Bureau of Economic Research.
- Moody's Investors Service. 2015. "Annual Default Study: Corporate Default and Recovery Rates, 1920-2014," Technical report.
- Nimark, K. 2014. "Man-Bites-Dog Business Cycles," *American Economic Review*, 104(8), 2320-2367.
- Reinhart, C., and K. Rogoff. 2020. "The Coronavirus Debt Threat," available at <https://www.wsj.com/articles/the-coronavirus-debt-threat-11585262515> (March 26, 2020).
- Reinhart, C., and K. Rogoff. 2009. *This Time Is Different: Eight Centuries of Financial Folly*, Princeton University Press.
- Stock, J.H. 2020, March. "Data Gaps and the Policy Response to the Novel Coronavirus," Working Paper 26902, National Bureau of Economic Research.
- Straub, L., and R. Ulbricht. 2013. "Credit Crunches, Information Failures, and the Persistence of Pessimism."
- Van Nieuwerburgh, S., and L. Veldkamp. 2006. "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53(4), 753-772.
- Velde, F.R. 2020. "What Happened to the U.S. Economy During the 1918 Influenza Pandemic? A View Through High-Frequency Data," SSRN Working Paper Series.
- Veldkamp, L., and J. Wolfers. 2007. "Aggregate Shocks or Aggregate Information? Costly Information and Business Cycle Comovement," *Journal of Monetary Economics*, 54, 37-55.
- Wang, H., Z. Wang, Y. Dong, R. Chang, C. Xu, X. Yu, S. Zhang, L. Tsamlag, M. Shang, J. Huang, et al. 2020. "Phase-Adjusted Estimation of the Number of Coronavirus Disease 2019 Cases in Wuhan, China," *Cell Discovery*, 6(1), 1-8.