



Epidemic responses under uncertainty

Michael Barnett^{a,1,2} , Greg Buchak^{b,1} , and Constantine Yannelis^{c,1}

Edited by Jose Scheinkman, Columbia University, New York, NY; received May 16, 2022; accepted November 20, 2022

We examine how policymakers react to a pandemic with uncertainty around key epidemiological and economic policy parameters by embedding a macroeconomic SIR model in a robust control framework. Uncertainty about disease virulence and severity leads to stricter and more persistent quarantines, while uncertainty about the economic costs of mitigation leads to less stringent quarantines. On net, an uncertainty-averse planner adopts stronger mitigation measures. Intuitively, the cost of underestimating the pandemic is out-of-control growth and permanent loss of life, while the cost of underestimating the economic consequences of quarantine is more transitory.

COVID-19 | ambiguity | model uncertainty | dynamic general equilibrium

The rapid spread of COVID-19 at the beginning of 2020 was accompanied by a vigorous debate over the costs and benefits of actions taken to mitigate the pandemic's spread. This debate occurred under significant uncertainty regarding key parameters relating to the threat posed by the virus, including death rates and infection rates, as well as uncertainty about the costs and effectiveness of mitigation efforts (1). Many policymakers, academics, and commentators in the media suggested that this uncertainty argued for a laxer quarantine and lockdown response. We examine this claim by formally exploring how optimal pandemic mitigation policies change when faced with significant uncertainty.

We account for two types of model uncertainty channels in our analysis. The first channel is uncertainty around the epidemiological model. We focus on two epidemiological parameters that characterize the severity of a contagious disease: the case fatality rate (CFR), or the fraction of individuals infected who die due to the disease, and the basic reproduction number \mathcal{R}_0 , or the number of people in an otherwise healthy population that a single disease carrier is expected to infect. Early estimates of the CFR ranged from a flu-like .08% to a catastrophic 13.04%. Estimating a CFR for a new disease while cases are ongoing is inherently difficult as cases must be closed through either recovery or death before a CFR can be computed (2). Similar difficulties* in estimating \mathcal{R}_0 led to estimates ranging from 1.5 to 12 (6). To highlight the wide variation in these estimates, Fig. 1 shows initial CFRs and \mathcal{R}_0 across many countries and US states.[†]

The second channel is uncertainty around the impacts of the policy response. In particular, we focus on the effectiveness and costs of mitigation policies. Experience revealed significant ex ante uncertainty in how effective various measures would be and how stringently households would follow them. Additionally, the unprecedented nature of global stay-at-home orders in a modern economic setting led to vast uncertainty about how they would limit productivity and disrupt supply chains. Acemoglu et al. (7), for example, state that “We hasten to emphasize that there is considerable uncertainty about... the exact economic damages caused by lockdowns (in part because neither the extent to which work from home can substitute for workplace interactions nor the knock-on effects of current measures on supply chains and worker-firm relations are yet well understood).”

Knowledge of both epidemiological parameters and the effectiveness and costs of mitigation policies informs the central trade-off of an optimal response: how to balance the public health benefits (and the resultant downstream economic benefits) of mitigation policies against their economic costs. Confronted with this uncertainty yet facing a concrete decision, policymakers and politicians speculated on the role that uncertainty

*Many academic papers, e.g., ref. 3, note that \mathcal{R}_0 is difficult to estimate because the provision of tests is not random but rather targets those showing symptoms or those thought to be at higher risk. Manski et al. (4) discuss in detail the wide range of estimates and highlight how a lack of testing and the presence of many asymptomatic carriers made measurement difficult. Atkeson et al. (5) provide a model explaining the range of estimates through behavioral responses.

[†]The insights of this model apply not only to new epidemic diseases but also to new strains of existing diseases. While we focus on exposition on the initial emergence of the new virus, throughout 2020 and 2021, new variants emerged. Like with the initial disease, there has been significant uncertainty about the death rates and the rate of transmission of these new variants, resulting in persistent uncertainty about the pandemic more generally. The introduction of vaccines led to a new source of uncertainty.

Significance

How forcefully should policymakers react to a pandemic given uncertainty about the infectivity and severity of a new disease or variant? Some policymakers argued that more uncertainty should lead to prioritizing the economic costs of lockdowns, rather than the public health costs of the disease's uncertain spread. In this paper, we embed a macroeconomic SIR model in a robust control framework and show that, on the contrary, uncertainty about the disease should lead to a more decisive response to the initial outbreak. Intuitively, uncertainty over the exponential growth rate of the disease means that it is optimal for policymakers to prioritize avoiding worst-case scenarios where the disease is highly contagious, spreading rapidly, and exceedingly lethal.

Author affiliations: ^aW. P. Carey School of Business, Arizona State University, Tempe, AZ 85287; ^bStanford University Graduate School of Business, Stanford, CA 94305; and ^cUniversity of Chicago Booth School of Business, Chicago, IL 60637

Author contributions: M.B., G.B., and C.Y. designed research; performed research; analyzed data; and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This open access article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

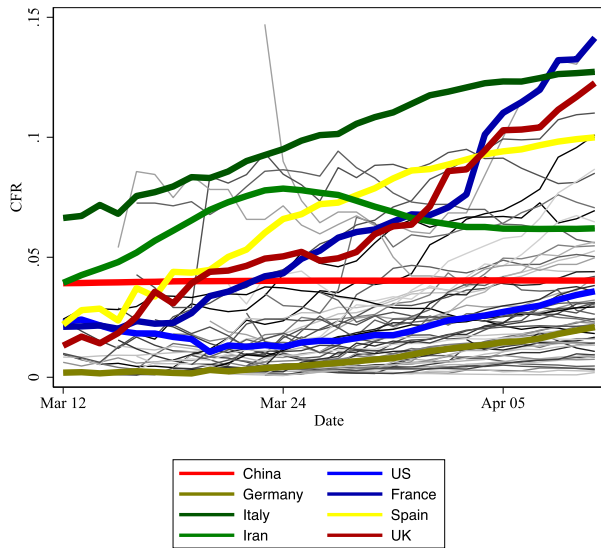
¹M.B., G.B., and C.Y. contributed equally to this work.

²To whom correspondence may be addressed. Email: michael.d.barnett@asu.edu.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2208111120/-DCSupplemental>.

Published January 6, 2023.

Panel A: Estimated CFR Rates by Country



Panel B: Estimated R_0 for Countries and States

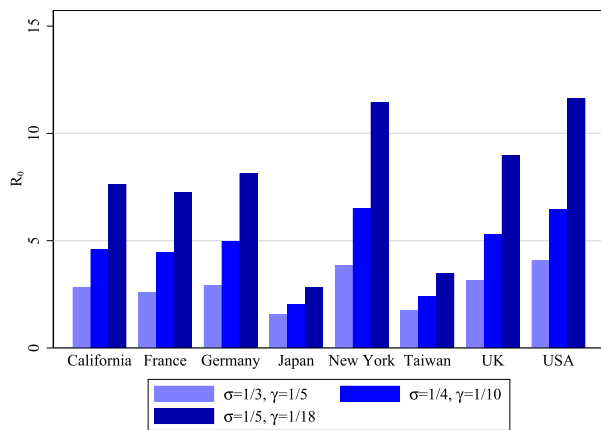


Fig. 1. Estimated CFR rates and R_0 . Panel A shows estimated CFRs for all countries with more than 1,000 cases and 100 deaths. Panel B shows estimated R_0 across countries and states given different incubation periods (σ) and durations (γ). Source: European Centre for Disease Prevention and Control and ref. 6.

should play in their decision-making: [New York Mayor De Blasio in a March 9 press conference](#) said “I am very resistant to take actions that we’re not certain would be helpful, but that would cause people to lose their livelihoods.” [Epidemiologist John Ioannidis remarked](#) “In the absence of data, prepare-for-the-worst reasoning leads to extreme measures of social distancing and lockdowns. Unfortunately, we do not know if these measures work.... This has been the perspective behind the different stance of the United Kingdom keeping schools open.”

Dynamic decision theory applied to economic models offers a rigorous way forward when confronted with parameter uncertainty. In particular, refs. 8 and 9, and more recently (10) suggest a max–min criteria whereby a policymaker selects the policy that would be optimal under a worst-case scenario. The worst-case scenario under consideration must be disciplined by what is reasonably consistent with the data. For example, an extremely contagious disease with an eventual 100% fatality rate is indeed a worst-case scenario but—parameter uncertainty notwithstanding—is not consistent with even the most pessimistic estimates. A CFR of 5%, however, while toward the upper end of estimates, may be a reasonable worst-case scenario to consider.

In our paper, we adopt a formalization of this idea through a smooth ambiguity approach, which provides a tractable framework to select a reasonable worst-case scenario. There are two key distinguishing features of the uncertainty-averse model. First, uncertainty aversion tilts the policymaker’s model weighting toward parameters with more substantial utility consequences. Therefore, the planner gives stronger consideration to worst-case scenarios, with the strongest tilts occurring for the most uncertain parameters. Second, this tilting evolves endogenously as the pandemic unfolds. The planner observes the state of the pandemic and, without learning about or knowing the true model parameters, adjusts the worst-case prior toward new potential worst-case outcomes based on how the pandemic has evolved over time.

We begin with a simple macroeconomic model with an epidemic without model uncertainty. The epidemic spreads according to a standard susceptible-infectious-recovered (SIR) process employed in epidemiology and more recently in macroeconomic models.[‡] Infected individuals are less productive, spread the disease, and may die. A policymaker can impose a quarantine of varying strictness. The quarantine slows the disease’s spread by removing individuals from the working population at the cost of temporarily reducing output. The optimal policy balances spread reduction against temporary economic losses and depends critically on the epidemiological parameters, the effectiveness of mitigation policies, and the cost of mitigation policies. We first show the wide range of optimal responses across different underlying parameters. These wildly varying comparative statics drive home the potential costs of model ambiguity yet by themselves do not offer a prescription for how a policymaker facing this uncertainty might actually respond.

To provide an answer, we adopt a smooth ambiguity approach to explicitly introduce parameter uncertainty into the policymaker’s decision problem. As a critical first step, we differentiate risk from uncertainty. Following refs. 11 and 12, risk refers to the range of possible outcomes in a model where the parameters are known. In contrast, uncertainty refers to the possibility that the model’s parameters are unknown or that the model itself is misspecified.[§] In our context, we introduce risk by allowing the disease to spread and kill nondeterministically. This risk gives rise to uncertainty by obscuring the true parameters governing the disease’s spread and lethality as well as the effectiveness and costs of mitigation policies. Facing this uncertainty, the policymaker considers possible outcomes across alternate parameter settings when making decisions. The set of parameters under consideration is disciplined by how far these parameters lie from the policymaker’s prior beliefs. We calibrate the model to match the US economy and explore how uncertainty influences optimal quarantine policy.

Our analysis shows that compared to a benchmark ambiguity-neutral planner, ambiguity aversion leads the optimizing planner to adopt a stronger and more persistent quarantine.

[‡]SIR models are standard tools in epidemiology used to model the spread of infectious diseases. The epidemiological SIR model computes the theoretical number of people infected with a contagious disease in a closed population over time. The models have three key elements: S is the number of susceptible, I is the number of infectious, and R is the number of recovered, deceased, or immune individuals. Recent studies in macroeconomics have incorporated SIR models into macroeconomics models. [Stanford Earth System Sciences notes](#) provide an introduction to the standard epidemiological SIR model.

[§]A large body of literature refers to Brownian shocks and time variation in exposure to Brownian shocks as uncertainty, for example, refs. 13 and 14. Our model does not directly include a behavioral private agent response, and instead, this is captured by the damage function, a modeling simplification we make for tractability. If agents respond to policy, this could either mitigate the economic damages or reduce the effectiveness of quarantine. Ref. 15 provides a more complex model in which households respond to policy.

An ambiguity-neutral planner with an equal-weighted prior across models⁴ adopts a relatively modest quarantine, with the quarantine isolating roughly 50% of the population around 8 wk into the pandemic and ending after roughly 15 wk. In contrast, an ambiguity-averse policymaker with the same prior immediately implements a quarantine policy of isolating 25% of the population, increasing it to 60% in 10 wk, and maintaining some level of quarantine for just under 25 wk. When simulating these policies under the true model parameters, which match the values implied by the prior, the ambiguity-averse policies lead to fewer deaths and a flatter infection curve.

Uncertainty about disease infectivity, fatality rates, and quarantine effectiveness pushes the planner to adopt more stringent quarantines, while uncertainty about the economic costs of quarantine pushes the planner to adopt less stringent measures. Our results make clear that the former are (endogenously) the planner's primary concern, while the latter are of secondary concern. Intuitively, underestimating the disease's severity or overestimating the quarantine's effectiveness can lead to the out-of-control spread of the virus and permanent pandemic deaths. In contrast, underestimating the economic consequences of quarantine leads to costs that are more transitory in nature. Thus, the planner is more willing to place greater probability weights on model distortions where these convex and permanent costs are potentially larger. While the model and probability distortions leading to these results can be quite substantial, we show that the distortions are statistically reasonable based on model detection error probabilities.

Our paper primarily links to literature on ambiguity and robust control beginning with refs. 8 and 9 and continued by refs. 10 and 16–20.[#] Recent work in finance and macroeconomics has emphasized the importance of different forms of ambiguity and uncertainty, for example, refs. 13, 14 and 24–29. Robust control methods have also been used to study the economic impacts of climate change, as in refs. 30–33, where climate damages can have both permanent growth effects and transitory level effects. However, in the pandemic context, policy decisions trade off temporary (through quarantine and temporary illness) and permanent (through death) implications based on how model uncertainty amplifies concerns about the worst-case outcome.

A key contribution of our paper is to introduce uncertainty to the discussion on economic responses to the COVID-19 epidemic. A number of studies have built macroeconomic frameworks, combining SIR models from epidemiology with macroeconomic models, such as refs. 15 and 35–40. These studies rely on calibrated parameters, which are often unknown. Parameter uncertainty is widely acknowledged by these studies, and authors typically use a range of values. For example, (7) note that: “We stress that there is much uncertainty about many of the key parameters for COVID-19 (4), and any optimal policy, whether uniform or not, will be highly sensitive to these parameters e.g., refs. 3, 41, and 42. So, our quantitative results are mainly illustrative and should be interpreted with caution”. Ref. 34 is one important paper that has considered the issues and importance of incorporating model uncertainty for decision makers confronting COVID-19. Our paper complements this

⁴We also explore three additional sets of prior beliefs over the models in *SI Appendix*: i) underestimating the pandemic, ii) correctly estimating the pandemic, and iii) overestimating the pandemic. These results reinforce the baseline model outcomes, while highlighting asymmetries that arise in these alternative scenarios that are relevant for policymakers to consider. In particular, ambiguity aversion is the most relevant when initially underestimating the pandemic because a larger fraction of the population gets the virus quickly and thus potential spread under the worst-case scenario is much faster.

[#]There is an important and extensive body of theoretical study on uncertainty of various forms dating to 11 and 21–23.




Fig. 2. Expected transition rates between states in the augmented SIR model.

work by solving a theoretical model that explicitly evaluates the implications of multiple sources of ambiguity.

The remainder of this paper is organized as follows. *A Simple Economic Model of an Epidemic* presents our model. *Model Solutions* describes how to account for uncertainty. *Numerical Results* presents simulation results and *Concluding Remarks* concludes.

1. A Simple Economic Model of an Epidemic

We introduce a simple economic model of an epidemic without model uncertainty before incorporating uncertainty in subsequent sections. Our model embeds a simple SIR framework into an economic model that allows us to speak to the costs of the disease as well as the costs and benefits of mitigation efforts.

A. Epidemic Model. A standard SIR model is characterized by three state variables: the number of susceptible individuals S_t , the number of infected individuals I_t , and the number of recovered individuals R_t . In addition, we include a state variable for the total population N_t to account for deaths from the pandemic. To simplify the model solution, we use in our analysis the SIR state variables defined as fractions of the total population, i.e., $s_t = \frac{S_t}{N_t}$, $i_t = \frac{I_t}{N_t}$, and $r_t = \frac{R_t}{N_t}$. Transitions between the different states in the model depend on β_t , the rate at which a susceptible becomes infected, ρ_t , the rate at which an infected recovers, and δ_t , the rate at which an infected dies. We have abstracted from births and deaths not related to the epidemic for simplicity, which can be easily incorporated into our framework but has no qualitative impact on our results.

Fig. 2 illustrates the transition rates between states in our model. Under our specification, the transition rates β_t , ρ_t , and δ_t are directly linked to the key structural disease parameters mentioned previously, the CFR and \mathcal{R}_0 , as follows:

$$\mathcal{R}_0 = \frac{\beta_t}{\gamma_t}, \quad CFR = \frac{\delta_t}{\gamma_t}, \quad \gamma_t = \rho_t + \delta_t.$$

We assume constant values, conditional on a given model, for the expected time of infection and rate of infection, i.e., $\gamma_t = \gamma$ and $\beta_t = \beta$. For the death rate of infected individuals δ_t , a critical issue that has been at the forefront during the COVID-19 pandemic is the fact that increased infections have led to an increased death rate due to limited resources for treatment resulting from increased hospitalizations. Similar to the frameworks used by refs. 40 and 39, we specify δ_t as an increasing function of i_t given by $\delta_t = \delta + \delta_+ i_t$. By definition, the recovery rate of infected individuals ρ_t will depend on this specification as well.

Departing from the standard model, we introduce Brownian shocks through W_t^\parallel . Two dimensions of W_t , which we denote W_i and W_d , are incorporated as parameter perturbations for β_t and δ_t with volatilities σ_i and σ_d , respectively. These shocks capture, for example, variability in exposure, comorbidities, mismeasurement, and random fluctuations in the number of susceptible, infected, recovered, and population size.

The state evolution equations we use in our analysis are given as follows:

$$\begin{aligned} ds_t &= -\beta s_t i_t dt + s_t i_t \delta_t dt \\ &\quad - s_t i_t \sigma_i dW_i + s_t i_t \sigma_d dW_d, \\ di_t &= \beta s_t i_t dt - \gamma i_t dt + i_t^2 \delta_t dt \\ &\quad + \sigma_i s_t i_t dW_i - i_t \sigma_d dW_d + i_t^2 \sigma_d dW_d, \\ r_t &= 1 - s_t - i_t, \\ \frac{dN_t}{N_t} &= -i_t \delta_t dt - i_t \sigma_d dW_d. \end{aligned}$$

While we specify these state evolution equations directly here, *SI Appendix* provides the evolution processes S_t , I_t , and R_t as well as the derivation of the evolution processes for s_t , i_t , and r_t by way of Ito's lemma.

A.1. Pandemic mitigation. We allow for pandemic mitigation through quarantine measures. Let q_t be the fraction of the population in quarantine at any period of time, where “quarantine” captures a wide range of policies such as school closures, business closures, and shelter-in-place orders. Quarantine prevents susceptible individuals from becoming infected. Given the mitigation policy q_t , the population laws of motion for the susceptible and infected become as follows:

$$\begin{aligned} ds_t &= -\beta s_t i_t (1 - \zeta q_t)^2 dt + s_t i_t \delta_t dt \\ &\quad - s_t i_t \sigma_i dW_i + s_t i_t \sigma_d dW_d, \\ di_t &= \beta s_t i_t (1 - \zeta q_t)^2 dt - \gamma i_t dt + i_t^2 \delta_t dt \\ &\quad + \sigma_i s_t i_t dW_i - i_t \sigma_d dW_d + i_t^2 \sigma_d dW_d. \end{aligned}$$

This specification mirrors that in ref. 40 in terms of the impact of the quarantine. $s_t(1 - \zeta q_t)$ and $i_t(1 - \zeta q_t)$ are the masses of susceptible and infected that meet. $\zeta \in [0, 1]$ captures the incomplete effectiveness of quarantine measures, e.g., meeting with family, shopping, or ignoring the policy altogether.

B. Economic and Public Health Model.

B.1. Preferences, production, and consumption. We focus our analysis on a social planner's problem, highlighting the optimal quarantine policy choice of a benevolent government or policy-maker who internalizes any externalities and seeks to maximize social welfare while confronting uncertainty. The planner has flow utility that depends on consumption C_t and a subjective discount rate κ and is given by**

$$U_t = \kappa \log C_t.$$

Log utility allows us to incorporate risk aversion in the simplest way into our framework, a relevant feature given the inclusion of Brownian shocks for our state variables. A linear

^{||}Formally, $W \doteq \{W_t : t \geq 0\}$ is a multidimensional Brownian motion where the corresponding filtration is denoted by $\mathcal{F} \doteq \{\mathcal{F}_t : t \geq 0\}$ and \mathcal{F}_t is generated by the Brownian motion between dates zero and t .

**We discuss the impact of including nonpecuniary losses for deaths from the pandemic as in refs. 15, 35, and 40 in *SI Appendix*. Such costs, and uncertainty about these costs, should serve to enhance the results we find in our main analysis.

production technology produces output Y_t with labor L_t and labor productivity A_t . Here, A_t includes the capital stock, which we hold fixed. Households consume everything that is produced so that

$$C_t = Y_t = A_t L_t.$$

The labor supply is determined by the total population, which varies with shocks and deaths from the pandemic, and the magnitude of the quarantine measures put in place to mitigate spread of the pandemic. The effective supply is therefore defined by:

$$\begin{aligned} L_t &= N_t [(1 - q_t)(s_t + \phi i_t + r_t)] \\ &= N_t [(1 - q_t)(1 - (1 - \phi)i_t)], \end{aligned}$$

where $s_t + \phi i_t + r_t$ is the effective available labor force, $\phi \in (0, 1)$ represents the amount by which an infected worker's productivity is reduced, and $1 - q_t$ is the nonquarantined fraction of the available labor force.^{††}

B.2. Productivity costs of mitigation. Beyond the clear costs of mitigation in reducing the available labor force, economic costs in the form of reduced productivity have also been of first-order concern. That is, even individuals who are not locked down may become less productive due to economic disruptions caused by lockdowns. To incorporate this additional channel, we assume that mitigation efforts can lead to economic costs in the form of reduced productivity. We model this formally by expanding our expression of productivity to follow a modified Ornstein–Uhlenbeck process as follows:

$$\begin{aligned} A_t &= \bar{A} \exp(z_t), \\ dz_t &= (-\alpha q_t - z_t) dt + \sigma_z dW_z, \end{aligned}$$

where W_z is the third dimension of the vector Brownian motion W_t . This process is mean-reverting, and so, shocks to productivity are transitory. Without quarantine, the long-term mean of z_t is zero. When quarantine measures are introduced, the new-long term mean is given by $-\alpha q_t$. The productivity impact persists while quarantine measures are in place, and as soon as they are ended, the process mean-reverts back to the long-term mean of zero. Thus, mitigation can have significant costs to economic productivity, but those costs are transitory as long as mitigation efforts are not permanent. This additional channel of economic costs to mitigation will interact with the existing forces of reduced labor force from mitigation and concerns about the costs of the pandemic, adding to the key trade-offs that the social planner must consider when determining optimal policy responses.

B.3. Arrival of a vaccine and cure. We assume that there is a constant arrival rate λ of a resolution of the epidemic arriving at some unknown time in the future T . Our specification, consistent with others in this area of research such as refs. 35 and 40, assumes that upon the realization of the resolution shock taking place, a cure and a vaccine are found so that all susceptible individuals are immune and all infected individuals recover. The arrival rate is set so that this resolution is expected to occur in about 1.0 y. The main impact of this assumption is that expectations about a resolution of the pandemic lead to amplification of the subjective discount rate of the planner, providing quantitatively more realistic policy responses. We provide the full details for the derivation of the model under this assumption in *SI Appendix*.

^{††}An alternate specification could target only infected or susceptible and infected workers for quarantine. Because most quarantine policies in practice have been untargeted, we adopt the untargeted specification. The results are easily extended to the targeted quarantine setting and are qualitatively similar.

C. Parameter Uncertainty. We motivated our paper with uncertainty over the pandemic model parameters as well as the quarantine policy model parameters. We incorporate uncertainty into our setting by assuming a discrete set Θ of possible models θ . Each $\theta \in \Theta$ corresponds to a set of parameters $\{\beta(\theta), \delta_t(\theta), \zeta(\theta), \alpha(\theta)\}$. The interpretation is that each model θ comes from an existing estimate for the true pandemic model and the true impacts of quarantine policy inferred from historical data, real-time information, or other sources. Each θ characterizes the state evolution equations as follows:

$$\begin{aligned} ds_t &= -\beta(\theta)s_t i_t(1 - \zeta(\theta)q_t(\theta))^2 dt + s_t i_t \delta_t(\theta) dt \\ &\quad - s_t i_t \sigma_i dW_i + s_t i_t \sigma_d dW_d, \\ di_t &= \beta(\theta)s_t i_t(1 - \zeta(\theta)q_t(\theta))^2 dt - \gamma i_t dt + i_t^2 \delta_t(\theta) dt \\ &\quad + \sigma_i s_t i_t dW_i - i_t \sigma_d dW_d + i_t^2 \sigma_d dW_d, \\ dz_t &= -\alpha(\theta)q_t(\theta) dt - z_t dt + \sigma_z dW_z, \\ \frac{dN_t}{N_t} &= -i_t \delta_t(\theta) dt - i_t \sigma_d dW_d. \end{aligned}$$

Crucially, the Brownian shocks we introduce into the dynamic evolution equations used in our analysis prevent the policymaker from immediately inferring the fundamental transition rates of the disease and the effectiveness and cost of quarantine measures as we explore the impacts of uncertainty and ambiguity in the model. How the planner confronts this uncertainty depends upon the decision framework implemented by the planner to determine social optimality. In what follows, we derive the model solutions for the planner's problem with and without aversion to model uncertainty as captured by the smooth ambiguity framework from the dynamic decision theory toolkit.

2. Model Solutions

We solve the social planner's problem with and without ambiguity-based model uncertainty. From this, we are able to make a direct comparison of the impact of uncertainty on the optimum quarantine choice to highlight the key forces behind the planner's uncertainty-adjusted policy decisions and the difference it has on the economic and pandemic outcomes in the model. In each case, the solution to the infinitely-lived social planner's problem is a recursive equilibrium defined by a socially optimal quarantine policy q_t^* that maximizes the social welfare or expected lifetime utility of the planner subject to the evolution of the stochastic process for the state variables s_t, i_t, N_t, z_t , as well as the pandemic and economic adding up constraints.^{‡‡} The equilibrium solution has a Markovian structure such that the value function that characterizes the solution, and the optimal quarantine policy are functions of the state variables s_t, i_t, N_t , and z_t . To derive the model's socially optimal outcomes, we solve for the social planner's value function from the Hamilton–Jacobi–Bellman (HJB) equation representing their optimization problem in a recursive format. First-order conditions characterizing the optimal policies are derived from this HJB equation and used to solve for the value function and the optimal quarantine choice.

We first solve the planner's problem conditional on a given model θ , deriving an optimal policy $q_t(\theta)$ without reference to ambiguity. The typical analysis would end here with a comparison of policies across $\theta \in \Theta$. This “outside-the-model uncertainty”

^{‡‡}We can omit the state variable r_t from our framework because the explicit adding-up constraint of the SIR model means that r_t is the residual after accounting for s_t, i_t , and N_t and is therefore redundant.

or sensitivity analysis exercise often undertaken in economics and other fields does not account for the social planner making decisions while confronting this uncertainty explicitly. In our numerical results, we show how disperse the optimal quarantine policy and pandemic outcomes can be across the set of models we consider.

The ambiguity-averse planner's solution incorporates concerns about uncertainty directly into the social planner's decision problem, building on the continuous-time smooth ambiguity framework developed in ref. 19, and applied in the analysis of ref. 32. The result is a minimum–maximum problem where the planner optimizes over constrained worst-case model distortions (minimization) and optimal mitigation policy (maximization). In contrast to a simple model averaging framework, this form of uncertainty for the decision maker incorporates the fact that the planner does not know what weights to place on the different potential models of the pandemic and explicitly incorporates this ambiguity into the planner's decision problem. The decision maker chooses an initial distribution of prior weights to place on the models and then distorts these baseline weights based on endogenously determined optimal adjustments arising from their aversion to uncertainty in the form of model ambiguity. As a result, uncertainty is explicitly incorporated into the planner's optimal policy choice through probability adjustments used to weigh the different models θ , providing an uncertainty-adjusted optimal policy that arises from the minimum–maximum problem optimization.

A. Optimal Policy Without Uncertainty. We first solve the social planner's problem conditional on a given model $\theta \in \Theta$. The social planner's problem is to maximize lifetime expected utility by choosing the optimal mitigation or quarantine policy $q_t(\theta)$. The planner's problem can be expressed as

$$\begin{aligned} V(s_t, i_t, N_t, z_t; \theta) &= \max_{q_t} E_0[e^{-\kappa(T-t)} \hat{V}(N_T, z_T) \\ &\quad + \int_0^T e^{-(\kappa+\lambda)t} \kappa \log C(q_t) dt | \theta], \end{aligned}$$

subject to economic and pandemic constraints. Note that $C(q_t) = A_t \times L(q_t)$ and $\hat{V}(N_T, z_T)$ is the continuation value postpandemic. We represent the social planner's problem using a Hamilton–Jacobi–Bellman (HJB) equation for the value function resulting from the social planner's optimization. There is an analytical solution for $\hat{V}(N_T, z_T)$ and an analytical simplification for the value function given by

$$V(s_t, i_t, N_t, z_t; \theta) = \log(\bar{A}) + \log N_t + \frac{\kappa}{\kappa + 1} z_t + v(s_t, i_t; \theta).$$

After incorporating these simplifications, the simplified PDE we solve for the planner's problem is given by

$$\begin{aligned} &(\kappa + \lambda)v(s_t, i_t) \\ &= \max_{q_t} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) \\ &\quad - \delta_t i_t - \frac{1}{2} i_t^2 \sigma_d^2 - \frac{\kappa}{\kappa + 1} \alpha q_t + v_{s_t} i_t \delta_t - v_{i_t} [\gamma - \delta_t i_t] \\ &\quad + v_i \beta s_t i_t (1 - \zeta q_t)^2 - v_s \beta s_t i_t (1 - \zeta q_t)^2 \\ &\quad + \frac{1}{2} [v_{ss} (\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2 + v_{ii} (\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)] \\ &\quad - v_{si} [\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t (1 - i_t) i_t^2], \end{aligned}$$

where we drop the θ notation for brevity. See *SI Appendix* for full details on the derivation and analytical simplification of the HJB equation. The optimal choice of mitigation q_t is the solution to a quadratic equation resulting from the first-order condition and is given by

$$q_t(\theta) = \frac{-B - \sqrt{B^2 - 4C}}{2},$$

$$B = -\frac{1 + \zeta(\theta)}{\zeta(\theta)} - \frac{\frac{\kappa}{\kappa+1}\alpha(\theta)}{2\beta(\theta)\zeta(\theta)^2(v_i - v_s)s_t i_t},$$

$$C = \frac{\kappa + \frac{\kappa}{\kappa+1}\alpha(\theta)}{2\beta(\theta)\zeta(\theta)^2(v_i - v_s)s_t i_t} + \frac{1}{\zeta(\theta)}.$$

The policy function depends on the parameters associated with θ , $\{\beta(\theta), \delta_t(\theta), \zeta(\theta), \alpha(\theta)\}$. We discuss these optimal policies without uncertainty in subsequent sections.

B. Optimal Policy with Uncertainty. To analyze the impact of model uncertainty, we implement a smooth ambiguity framework.^{§§} To do this, we return to the full discrete set Θ of possible models θ for the pandemic as noted above. We first specify prior probability weights for the set of models $\theta \in \Theta$, by assigning a probability weight $\pi(\theta)$ to each model θ , satisfying

$$\pi(\theta) \geq 0 \quad \forall \theta \in \Theta, \quad \sum_{\theta \in \Theta} \pi(\theta) = 1.$$

Like the alternative models in our set, the prior probability weights are assumed to come from historical data or real-time observational inference.

We then allow for uncertainty aversion by using a penalization framework based on conditional relative entropy. This framework allows the planner to consider alternative distributions or sets of weights $\tilde{\pi}(\theta)$ across the set of conditional models in a way that is statistically and quantitatively reasonable in order to derive optimal policy that is robust to possible worst-case models that the planner is concerned with due to the model uncertainty. In essence, the social planner considers worst-case distributions but penalizes probability weights that are far from the planner's prior beliefs. This penalization is based on an ambiguity aversion parameter, which penalizes the magnitude of the deviation of the distorted probability weights from the prior weights. The parameter ξ_a is the ambiguity parameter that determines the magnitude of this penalization, calibrated to ensure that worst-case models considered by our decision maker are statistically and quantitatively reasonable. Large values of ξ_a imply low aversion to ambiguity, while small values of ξ_a imply strong aversion to ambiguity. Relative entropy, defined as the expected value of the log-likelihood ratio between two models or the expected value of the log of the Radon–Nikodym derivative between models, is the measure used to determine the magnitude of the deviation of the distorted probability weights from the prior weights. See ref. 19 for details about relative entropy in this setting. Using relative

^{§§}The most common alternative to a smooth ambiguity approach is a robust preference approach as in ref. 16. A key advantage of our approach is that the uncertainty here is structured into alternative models as characterized by explicit sets of key parameters and distorted probabilities. This type of structured uncertainty analysis allows us to examine explicitly how prior model weights are distorted and therefore determine which models are of most interest to the uncertain planner when making optimal policy choices that are robust to the existing ambiguity. In *SI Appendix*, we provide an extension of the model where we apply the continuous-time robustness framework studied in refs. 10, 16, 17, and others. While we are able to demonstrate the theoretical differences in how the different types of decision theoretic frameworks impact the structure of the planner's problem, there are potentially interesting quantitative differences that we leave for future research.

entropy means that we are considering only relatively small distortions from the baseline model, but even small distortions can have significant impacts on optimal policy. To give a concrete example in the context of COVID-19, it may be relatively easy to observe the number of people who died from the pandemic but difficult to observe the number of people who were infected. On the basis of this data, it is difficult to tell whether the disease has a very high spread rate (\mathcal{R}_0) and a low death rate (CFR), or a low spread rate and a very high death rate, yet the optimal response is likely to be very different under these scenarios.

While we have incorporated additional structure and complexity to the model to account for model uncertainty, the resulting household or social planner problem remains tractable and similar to the previous, no uncertainty problem, and is given by

$$V(s_t, i_t, N_t, z_t)$$

$$= \max_{q_t} \min_{\tilde{\pi}(\theta)} E_0 \left[e^{-\kappa(T-t)} \hat{V}(N_T, z_T) + \int_0^T \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \right.$$

$$\left. \times \left\{ e^{-(\kappa+\lambda)t} \left(\kappa \log C(q_t) + \xi_a \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \right) \right\} dt \right],$$

subject to the economic and pandemic constraints and the dynamics of the state variables relevant to the model. Note that the planner problem now has two layers of expectation and optimization. The outer expectation is over unknown outcomes of the Brownian shocks, while the inner expectation is over the possible models and is denoted by the sum over $\theta \in \Theta$ for probabilities $\tilde{\pi}(\theta)$. The inner minimization represents the planner considering the worst-case outcomes across possible models given the policy q_t , while the outer maximization represents the planner choosing the optimal quarantine policy q_t , understanding that it will be evaluated against the worst-case probability distribution. The term $\xi_a \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)}$ is the relative entropy penalization term with uncertainty parameter ξ_a , prior probability $\pi(\theta)$, and distorted probability $\tilde{\pi}(\theta)$.

As before, the social planner's solution is characterized by a recursive Markov equilibrium for which an equilibrium solution is defined as before. The HJB equation resulting from this modified household or social planner optimization problem which characterizes the socially optimal solution is now given by

$$(\kappa + \lambda)v(s_t, i_t)$$

$$= \max_{q_t} \min_{\tilde{\pi}(\theta)} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t)$$

$$+ \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \left\{ -\frac{\alpha\kappa}{\kappa + 1} q_t + \delta_t [-i_t + (v_i i_t + v_s s_t) i_t] \right.$$

$$+ (v_i - v_s) \beta s_t i_t (1 - \zeta q_t)^2 - v_i i_t \gamma \left. \right\} - \frac{1}{2} i_t^2 \sigma_d^2$$

$$+ \frac{1}{2} [v_{ss} (\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2 + v_{ii} (\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)]$$

$$- v_{si} [\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t (1 - i_t) i_t^2] + \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)},$$

where again, except for the π 's and $\tilde{\pi}$'s, we have suppressed the notation for θ for brevity. Taking the first-order condition for this problem, and imposing $\sum \tilde{\pi}(\theta) = 1$, we find that the optimally distorted probability weights are given by

$$\tilde{\pi}(s_t, i_t; \theta) \propto \pi(\theta) \exp\left(-\frac{1}{\xi_a} \left\{ -\frac{\alpha(\theta)\kappa}{\kappa + 1} q_t \right. \right.$$

$$\begin{aligned}
& + \delta_t(\theta)(-i_t + v_{s_t}i_t + v_i i_t^2) \\
& + \beta(\theta)s_t i_t(1 - \zeta(\theta)q_t)^2(v_i - v_s)).
\end{aligned}$$

As the $\tilde{\pi}(s_t, i_t; \theta)$ in the model are optimally determined and state dependent, the magnitude of the ambiguity considered by the social planner when making optimal policy decisions will depend on the current state of the pandemic and evolves dynamically.

The optimal choice of mitigation q_t has a similar functional form, given by

$$\begin{aligned}
q_t &= \frac{-B - \sqrt{B^2 - 4C}}{2}, \\
B &= -\frac{\tilde{\beta}\zeta + \tilde{\beta}\zeta^2}{\tilde{\beta}\zeta} - \frac{\tilde{\alpha}\kappa}{\kappa+1} \frac{1}{2\tilde{\beta}\zeta^2(v_i - v_s)s_t i_t}, \\
C &= \frac{\kappa + \frac{\tilde{\alpha}\kappa}{\kappa+1}}{2\tilde{\beta}\zeta^2(v_i - v_s)s_t i_t} + \frac{\tilde{\beta}\zeta}{\tilde{\beta}\zeta^2},
\end{aligned}$$

where the terms $\tilde{\beta}\zeta$, $\tilde{\beta}\zeta^2$, and $\tilde{\alpha}$ are given as before by

$$\begin{aligned}
\tilde{\beta}\zeta &= \sum_{\theta} \tilde{\pi}(s_t, i_t; \theta)\beta(\theta)\zeta(\theta), \\
\tilde{\beta}\zeta^2 &= \sum_{\theta} \tilde{\pi}(s_t, i_t; \theta)\beta(\theta)\zeta(\theta)^2, \\
\tilde{\alpha} &= \sum_{\theta} \tilde{\pi}(s_t, i_t; \theta)\alpha(\theta).
\end{aligned}$$

Now that ambiguity is incorporated into the optimal policy choice, the planner tilts the probability weights toward certain models based on potential worst-case outcomes so that the planner can respond in a robustly optimal way in the face of uncertainty. Importantly, the policy function depends on the parameters associated with the various θ models. But rather than solving for an optimal policy for each model and choosing either our preferred specification or taking a weighted average across model solutions based on a prior weighting, the optimally uncertainty-adjusted parameters are incorporated directly into the solution for the value function and the optimal policy choice.

We note that this analysis abstracts from any form of Bayesian learning. While learning is certainly an interesting consideration when thinking about the planner's optimal response to a pandemic, we find our setting valuable to consider for a number of reasons. First, the tractability of the smooth ambiguity framework in our analysis is particularly valuable for providing intuition about the implications of uncertainty. The characterization of ambiguity is condensed to a single dimensional parameter for uncertainty aversion rather than the potentially high-dimensional complexities or additional state variables than can arise from models of learning. Second, the rapid development of the COVID-19 pandemic and extreme difficulty in determining the true model for policymakers responding in real-time based on imperfect data, numerous virus variants, and an incomplete understanding of the effectiveness of certain factors influencing infections and deaths make this assumption a reasonable and realistic starting point. Recent work by Baek et al. (43) shows that, even with fixed parameter values, one can only hope to predict SIR model outcomes once a pandemic has nearly reached its peak infection rate due to limited testing and asymptomatic individuals. Finally, the arrival of new variants after the initial spread, each with unknown infectiousness and lethality, highlights the point that

uncertainty over the parameters of the disease overall may remain even as epidemiologists learn about individual variants.

3. Numerical Results

We now provide numerical results from simulations based on the theoretical solutions provided. Calibration of parameters is discussed first, followed by solutions for the models without uncertainty or "outside-the-model uncertainty" sensitivity analysis, and then by solutions for the model with uncertainty or "inside-the-model uncertainty" sensitivity analysis.

A. Parameter Values. This section discusses the parameter values used in the main calibration. These parameters are shown in the table below, and we discuss the parameter choices now. For the economic side of the model, we normalize the working population so that $L_0 = 1$ and set initial productivity to $A_0 = 20/12 = 1.667$ so that output in the nonpandemic version of the model ($A_0 \times L_0$) matches recent, prepandemic data on US GDP of \$20 trillion dollars annually or \$1.667 trillion dollars monthly. We choose an annual discount rate of 3%, and so the subjective discount rate κ is given by $\kappa = 0.0025$ for the baseline analysis. We assume that the expected arrival time for a vaccine is 1.0 y, so that the value of the arrival rate is given by $\lambda = \frac{1}{1.0} \times \frac{1}{12} = 0.083$.

For values of the mitigation effectiveness ζ and productivity costs of mitigation α parameters, we rely on estimates from the literature. Yamamoto et al. (44) estimate the effectiveness of quarantine measures for US states over time using an SIRD model and data on infections and stay-at-home orders. We choose our set of ζ values to be $\{0.35, 0.65\}$, which are close to their estimated upper and lower bounds. Barrot et al. (45) estimate the GDP losses for numerous European countries based on 6-wk of social distancing measures. We choose values for α in order to match GDP impacts of mitigation close to the average ($\approx 6\%$) and near the upper end ($\approx 8\%$) of the (45) estimates based on a fixed value of $q_t = 0.5$ over 6 wk. The GDP impact of social distancing measures in our model when there are no productivity impacts, and so $\alpha = 0$, is 6.25%. The GDP impact of social distancing measures in our model when $\alpha = 0.8$ is approximately 9.0%. Thus, we choose our set of values for α to be $\{0, 0.8\}$.

For the pandemic model parameters, we use values from various studies (including refs. 6, 41, 46, and 35, and estimates from the European Centre for Disease Prevention and Control) to set the expected time infected γ , the case fatality rate CFR, and basic reproduction number \mathcal{R}_0 , which allows us to pin down the infection rate β , the death rate δ , and the recovery rate ρ .

Table 1. Parameter values

Parameter	Variable	Value
Subjective discount rate	κ	0.0025
Nonpandemic output	$A \times \bar{L}$	1.667
Productivity costs of mitigation	α	$\{0, 0.8\}$
Infection severity	ϕ	0.4
Quarantine effectiveness	ζ	$\{0.35, 0.65\}$
Arrival rate of vaccine	λ	0.0833
Reproduction number	\mathcal{R}_0	$\{1.5, 4.5\}$
Initial case fatality rate	CFR	$\{0.005, 0.020\}$
Infection half-life	γ	30/18
Death rate convexity	δ_+	$5 \times \frac{CFR}{\gamma}$
Volatility	$\sigma_i, \sigma_d, \sigma_z$	$\{0.075, 0.030, 0.005\}$
Ambiguity parameter	ξ_a	$\{\infty, 0.0032\}$

The value of γ is held fixed at $\gamma = \frac{30}{18}$ or an expected duration of infection of 2.5 wk. The set of underlying models used in our analysis use initial values of CFR in the set $\{0.005, 0.02\}$ and values of \mathcal{R}_0 in the set $\{1.5, 4.5\}$. For the state-dependent death rate, rather than adding an additional set of parameters, we choose to scale the initial CFR values so that $\delta_+ = 5 \times \delta \times i_t$. This means that if 20% of the population were to become infected, i.e., if $i_t = 0.2$, then the CFR would be double the initial value. These values are well within the range of values across these different studies. The value of ϕ is 0.4, consistent with estimated values for the fraction of infected individuals who are asymptomatic given by the [CDC COVID-19 Pandemic Planning Scenarios website](#).

For σ_i , which is the volatility associated with β , we use the estimated reproductive number values \mathcal{R}_t of ref. 47 available through the [Tracking R website](#). To calculate a monthly time series standard deviation. This is done by taking the standard deviation calculated from the first 7 d of the COVID-19 outbreak for countries with outbreaks that occurred before global lockdown orders were implemented in order to capture the value for \mathcal{R}_0 , converting those values to a monthly frequency, and then averaging across these values to get a global value.^{¶¶} For the volatility σ_d , which is the volatility associated with δ_t , we use data from the Center for Systems Science and Engineering in the Whiting School of Engineering at Johns Hopkins University available through the [CSSE GitHub repo](#). To derive the global CFR, calculate the time series standard deviation using daily data from March 1, 2020, to March 31, 2020, and then convert this value to a monthly frequency.^{##} For σ_z , we use a value that matches that of ref. 32, in which the authors calculate a time series standard deviation of output from aggregate data.

Finally, we must also specify values for the uncertainty parameter in our model ξ_a for smooth ambiguity. Our main value of the uncertainty parameter, $\xi_a = 0.0032$, imposes what we view and aim to show is a reasonable amount of uncertainty aversion to demonstrate the potential magnitude of uncertainty impacts. This value can be difficult to interpret on its own and is best interpreted by way of model detection error probabilities implied by these parameter choices and the distorted model probabilities provided in the analysis. Furthermore, anecdotal evidence on model spreads implied by the recent estimates of COVID-19 parameter values which guide the values we use verify that our distorted values remain within a reasonable region for the given choice of ξ_a .

B. Outside-the-Model Uncertainty: Sensitivity Analysis. We first provide simulated outcomes of the model based on different pandemic models without the planner accounting for uncertainty in their optimal decision. This corresponds to what is typically termed a sensitivity analysis and illustrates the wide range of optimal responses that depend on the underlying model parameters. Fig. 3 shows the spread of outcomes for d_t and q_t across the different model cases. The spreads are across all model outcomes for $\mathcal{R}_0 \in \{1.5, 4.5\}$, initial CFR $\in \{0.005, 0.02\}$, quarantine effectiveness $\zeta \in \{0.35, 0.65\}$, and quarantine productivity cost parameter $\alpha \in \{0, 0.8\}$.

^{¶¶} We further constrain the list of countries to only those with daily volatility less than 0.25 to remove outliers and those with \mathcal{R}_0 to remove countries who implemented preemptive lock-down orders

^{##} While there is likely to be substantial measurement error impacting these volatility calibration calculations, the values of σ_i and σ_d have no qualitative impact and essentially no quantitative impact on q_t . The values do matter when we calculate the detection error probabilities, as they determine how large of a model distortion can be masked by the Brownian shocks. We therefore choose values of σ_i and σ_d that are conservative in order to avoid overstating what is a statistically reasonable amount of uncertainty to consider based on our choice of ξ_a .

Fig. 3 indicates very different policies q and outcomes d depending on parameters. Generally speaking, a higher reproduction rate, a higher death rate, a lower cost of quarantine, and a higher effectiveness of quarantine measures lead to more quarantine, all else being equal. Thus, we can anticipate that these models represent the worst-case from the planner's perspective when making optimal policy choices. In terms of the differences in policy choices across models, optimal quarantine policies are anywhere between no quarantine at all up to 80% of the population under quarantine, with the duration of policy measures lasting between 0 wk and approaching 1 y. The fraction of death in the population resulting from these policies varies by an order of magnitude, running between nearly 0.0025% and almost 4%. Observe that these quarantine choices and resulting death rates are made by a policy maker who knows the true parameters and is reacting optimally, and in that sense, they are best-case outcomes under each scenario. However, the dramatic

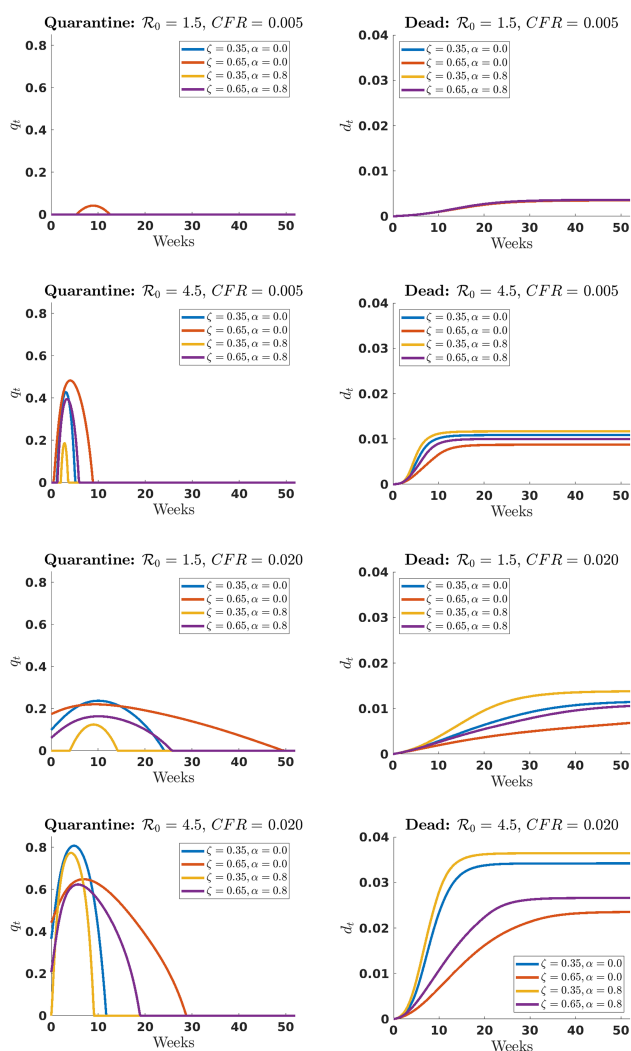


Fig. 3. Outside the model uncertainty, Notes: These figures show the range of outcomes and policy responses across 16 potential models of the pandemic that vary by \mathcal{R}_0 , CFR, α , and ζ . The *Left* column shows optimal quarantine policies by model, and the *Right* column shows the fraction of the population that dies by model. The *Top row* shows model results where $\mathcal{R}_0 = 1.5$, CFR = 0.005, $\alpha \in \{0.0, 0.8\}$, and $\zeta \in \{0.35, 0.65\}$. The *second row* shows model results where $\mathcal{R}_0 = 4.5$, CFR = 0.005, $\alpha \in \{0.0, 0.8\}$, and $\zeta \in \{0.35, 0.65\}$. The *third row* shows model results where $\mathcal{R}_0 = 1.5$, CFR = 0.02, $\alpha \in \{0.0, 0.8\}$, and $\zeta \in \{0.35, 0.65\}$. The *Bottom row* shows model results where $\mathcal{R}_0 = 4.5$, CFR = 0.02, $\alpha \in \{0.0, 0.8\}$, and $\zeta \in \{0.35, 0.65\}$.

variation in the magnitude and duration of quarantine policies suggests that an inappropriate response could lead to even more dramatically different outcomes. These differences highlight the likely significant role that accounting for model uncertainty will play in determining an optimal quarantine policy when the social planner incorporates ambiguity aversion.

C. Inside-the-Model Uncertainty: Smooth Ambiguity. The previous section highlighted drastically different responses and outcomes given different parameters as well as the likely worst-case outcomes for the model from the social planner's perspective. In this section, we examine how a policymaker explicitly accounting for these differences might respond. We assume that the true values match the simple averages of the possible parameters, with $\mathcal{R}_0 = 3.0$, initial CFR = 0.0109, costs of quarantine measures $\alpha = 0.4$, and quarantine effectiveness $\zeta = 0.5$. We then examine the optimal quarantine and resulting deaths from an "uncertainty neutral" solution where the policymaker simply applies the assumed prior weights to the models and the "uncertainty averse" policy where the planner starts from the same prior parameters but adjusts these weights due to concerns about ambiguity to derive the optimal solution. The assumed prior belief we consider in our main results is for the case where the planner gives equal weight to each possible model. In *SI Appendix*, we compare three additional sets of prior beliefs over the models: i) underestimating the pandemic, ii) split-estimating the pandemic, and iii) overestimating the pandemic. Each case highlights important scenarios the planner may face and how uncertainty influences policy choices and model outcomes in those different scenarios.

Fig. 4 shows the infected, dead, and quarantine outcomes for the uncertainty neutral and uncertainty averse cases of the equal-weighted prior scenario. Fig. 5 shows how the uncertainty averse planner's prior probabilities over each model are adjusted and distorted over time. Fig. 6 shows how the planner's beliefs over \mathcal{R}_0 , the CFR, the mitigation policy costs α , and the mitigation policy effectiveness ζ evolve correspondingly.

Starting with the uncertainty neutral results in red in Fig. 4, we see that quarantine policy starts at 0, then quickly moves up to around 50% near 8 wk, and then quickly drops back to 0 by 15 wk. The underlying pandemic peaks at around 15%

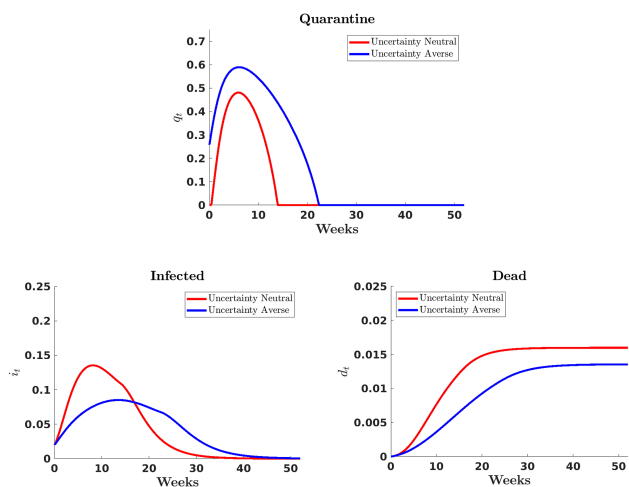


Fig. 4. Outcomes with and without uncertainty, equal weighted prior. Notes: These figures show the fraction of the population quarantined (*Top*), infected (*Bottom Left*), and dead (*Bottom Right*) under (i) the uncertainty neutral model response (red) and (ii) the uncertainty averse model response (blue).

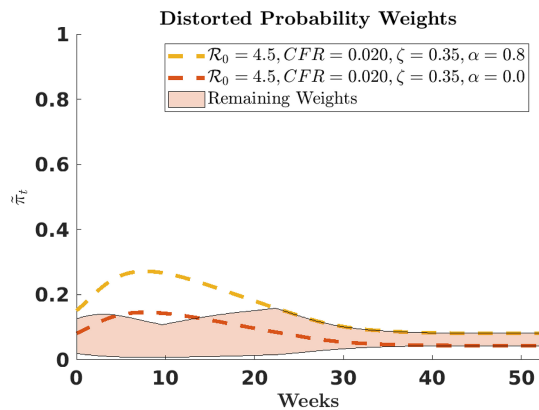


Fig. 5. Uncertainty aversion distorted model probabilities. Notes: This figure shows the distorted probability weights for the uncertainty averse planner. The dashed lines are the probability weights on the most severe epidemiological scenario ($CFR = 0.02$, $\mathcal{R}_0 = 4.5$) with the least effective pandemic mitigation scenario ($\zeta = 0.35$). The red dashed line is the model with the lowest economic cost of quarantine ($\alpha = 0.0$), and the yellow dashed line is the model with the highest economic cost of quarantine ($\alpha = 0.8$). The red-shaded area shows the range of distorted probability weights for the remaining models.

infected near 8 wk and results in about 1.5% dead overall. The impact of uncertainty aversion in blue is an emphatic increase in quarantine measures in terms of initial mitigation (around 25%), overall magnitude (above 60%), and in persistence of mitigation measures (just under 25 wk). The resulting impact on the evolution of the pandemic is that infections never go above 10%, and deaths are reduced by between a quarter to a half a percentage point.

These results demonstrate clearly that the concerns related to uncertainty that prompt the planner to implement stronger mitigation (lower quarantine effectiveness, higher death, and infection rates) dominate the concerns related to uncertainty that motivate the planner to reduce mitigation efforts (concerns about the cost of mitigation measures). Examining the distorted probabilities in Fig. 5 and the distorted model parameters of the uncertainty-averse planner in Fig. 6 highlights why this is the case. The two models in Fig. 5 that receive the largest increase in their prior probability weight are precisely those models with high \mathcal{R}_0 , high CFR, and low mitigation effectiveness ζ (yellow and red dashed lines). And while the model of these two that also assumes the highest economic costs of quarantine (yellow dashed line) receives the highest weight distortion, and is therefore the clear worst-case model in the planner's mind, the model from these two with the lowest economic costs of quarantine (red dashed line) is not too far behind. This result suggests that the economic costs are secondary to the other worst-case model concerns. It is also clear that the effects of model uncertainty are dynamically evolving and amplified as infections and deaths increase. The remaining model probabilities (red-shaded region) remain relatively low throughout, with some dropping to allow for a shift in probability to the worst-case models.

Fig. 6 shows the distorted model parameters implied by the planner's aversion to model ambiguity (blue lines) compared with the true model parameters (red lines). First, the uncertainty-averse planner responds as if the infection rate is persistently higher, with the distortion amplified as infections spike, and deaths increase. The dynamics of the death rate distortion are similarly influenced by the severity of the pandemic, though the spike in CFR is not as high as the uncertainty neutral case where the elevated number of infections has a more pronounced effect. The

quarantine effectiveness parameter is distorted downward, with the effect again amplified by the pandemic severity. However, this distortion quickly dissipates once the planner stops implementing quarantine measures around 25 wk. The distortions for the infection and death rates also distinctly drop off at this same time when mitigation measures are stopped and infections are quickly declining; however, these distortions persist longer than the quarantine effectiveness parameter distortion. The parameter for the economic costs of quarantine measures is persistently distorted up but lacks dynamic variation as the pandemic progresses. These results reemphasize what was highlighted by the distorted probability weights shown in Fig. 5, that the dominating concerns about model uncertainty, and therefore the model dimensions where the most probability weight is shifted to create the most prominent model distortions as the epidemic evolves, are the infection and death rates, followed by the quarantine effectiveness, and then the costs of quarantine measures. Because the equal-weighted prior on the possible models used in this setting avoids persistent biases created by alternative assumed priors, such as those used in the cases examined in the appendix, the effects of model uncertainty are demonstrated quite sharply here.

A useful diagnostic is to examine whether the optimal probability distortions are reasonable in the sense that they could not be rejected by observed data. We utilize the tool of detection error probabilities, shown in Fig. 7, for this. These probabilities, now commonly used in the literature on model uncertainty, are based on the model discrimination methods proposed by Chernoff (48) and Newman et al. (49). The bound quantifies the probability of making a type I or type II error when choosing between two models where the prior between the two models is 50–50. Following ref. 16, the detection error probabilities are calculated using a log-likelihood ratio-based model selection criterion that compares model outcomes for repeated simulations using the baseline and worst-case models. Under a standard heuristic, a value below 10% or 5% would indicate that a policymaker should reject the distorted model in the light of data generated by the baseline model. Fig. 7 shows clear dynamic shifts in the detection error probabilities for the uncertainty averse, with the probability starting just above 10% and then converging toward 50% after

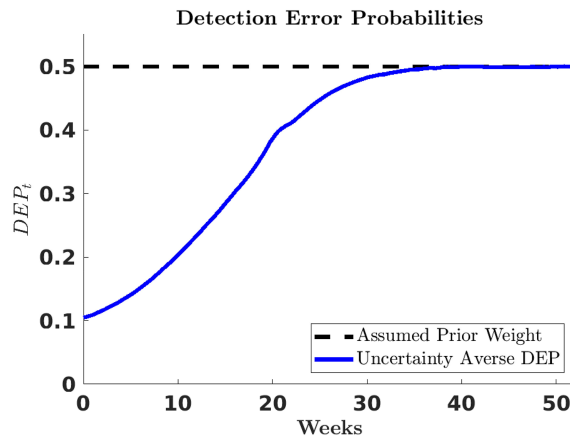


Fig. 7. Detection error probability comparison, Notes: The figure shows the detection error probability for model discrimination between the worst-case and baseline models in the uncertainty-averse (blue line) setting as well as the assumed prior weight on the baseline and worst-case models (dashed black line).

52 wk as infections converge toward zero. At that point, when the pandemic has largely run its course and the distorted weights are converging back toward the equal-weighted prior, the two models are nearly indistinguishable and the weighting on the worst-case and true models is close to the original 50% values. We find that the detection error probability never drops below a standard statistical significance value of 10%, let alone the often-used 5% or 1% values, and thus we interpret this magnitude of uncertainty distortion as statistically reasonable.

The central driving force in determining optimal policy choices under uncertainty aversion in our model setting is the trade-off between permanent and transitory costs in the model. The costs of a more severe pandemic are additional deaths as well as additional infections that can lead to more deaths, which has a permanent negative economic cost, whereas the economic costs of quarantine measures in the model are transitory. As a result, the uncertainty averse social planner places more emphasis on the worst-case epidemiological models which have the highest permanent costs, while still accounting to some degree for the transitory economic costs of increased quarantine measures. The result of the planner's optimal uncertainty adjustment is quarantine policy that is amplified and more persistent than in the uncertainty neutral setting in order to limit the number of infections and minimize the permanent cost of deaths. This amplification effect of uncertainty aversion on optimal quarantine measures is restrained by increasing marginal costs of additional quarantine measures, as well as the uncertainty about how severe those costs could be. Eventually, as the pandemic winds down enough that the planner no longer entertains the most severe potential worst-case epidemiological outcomes that seemed plausible during the peak of the pandemic, enhanced quarantine measures drop as well.

4. Concluding Remarks

In this paper, we embrace the admonition of ref. 50, that “[s]ince all models are wrong, the scientist must be alert to what is importantly wrong,” by studying optimal quarantine policy when allowing for uncertainty in models of epidemics or pandemics. Our main results focus on the role of uncertainty aversion in a smooth ambiguity-based decision problem. With new diseases, or diseases that have had only small outbreaks, there is often

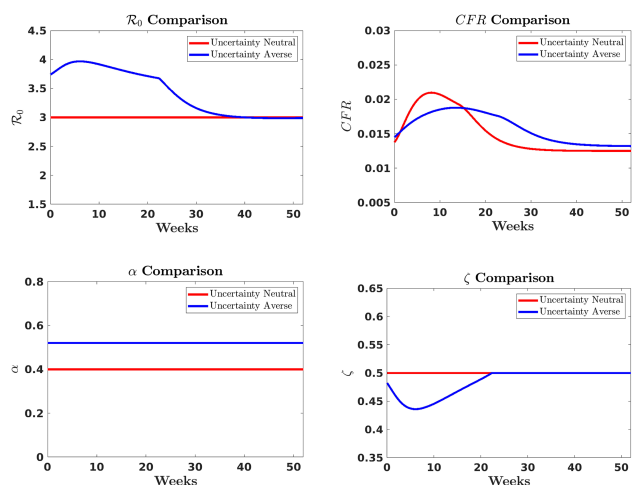


Fig. 6. Distorted model parameters: \mathcal{R}_0 , CFR , α , ζ . Notes: The figures show the implied \mathcal{R}_0 (Top Left), CFR (Top Right), α (Bottom Left), and ζ (Bottom Right) for the uncertainty neutral case (red lines) and the uncertainty averse case (blue lines). Each case uses the simulation path resulting from its own optimal policy.

significant uncertainty about key parameters which determine the overall consequences of an epidemic. The uncertainty-averse planner tilts the prior toward deviations with more substantial utility consequences, and this tilting evolves endogenously as the pandemic unfolds. The planner consequently emphasizes the worst-case scenarios in their decision problem by using the distorted prior to determine the optimal quarantine response when confronting aversion to uncertainty about the possible realization of worst-case outcomes.

The numerical results highlight how the planner trades off uncertainty concerns from epidemic and economic channels. Concerns about higher infection rates, higher fatality rates, and decreased effectiveness of quarantine measures increase the planner's motivation to implement quarantine measures, whereas uncertainty about the economic costs of quarantine pushes the planner to reduce quarantine measures. Taken together, the uncertainty-averse planner pushes for quarantine measures that are higher initially, peak higher, and persist longer than the planner who is uncertainty neutral. These effects are fairly substantial, even for levels of uncertainty aversion that are statistically reasonable based on detection error probabilities as a model selection criterion. The key to these results is that the planner emphasizes the permanent negative costs of the pandemic that are manifested as increased deaths over the more transitory economic costs of increased quarantine measures.

Our analysis provides a framework under which uncertainty and model misspecification can be incorporated into

macroeconomic models of epidemics. Our work emphasizes that uncertainty can play a large role in determining the optimal policy response to a new disease. Economists and epidemiologists, rather than using a range of parameters, can use our framework to explicitly model uncertainty. Future work can focus on making these models more tractable for policymakers, who often have to make decisions in real time, and further examine how economic agents modify their own behavior in response to uncertainty-averse policymakers.

Data, Materials, and Software Availability. All study data are included in the article, *SI Appendix*, and/or the accompanying online repository at <https://github.com/mbarnet0/Covid>.

ACKNOWLEDGMENTS. We are grateful to Scott Baker, Nick Bloom, Buz Brock, Lars Hansen, Zhiguo He, Murray Frank, Steve Kaplan, Ralph Koijen, Seth Pruitt, Laura Veldkamp, and Tom Sargent for helpful comments and discussions as well as seminar participants at the University of Arizona, Stanford, the University of Chicago, Arizona State University, the Midwest Finance Association Annual Meeting, and the Blue Collar Working Group at the University of Chicago. We are grateful to Livia Amato, Joseph Huang, Greg Tracy, and Minghao Yang for superb research assistance. Finally, our paper benefited tremendously from the comments and suggestions from the editor and two anonymous referees. Constantine Yannelis is grateful to the Fama Miller Center for generous financial support. Part of this research was performed while Michael Barnett was visiting the Institute for Mathematical and Statistical Innovation (IMSI), which is supported by the National Science Foundation (Grant No. DMS-1929348).

1. N. Chater, Facing up to the uncertainties of Covid-19. *Nat. Hum. Behav.* **4**, 439 (2020).
2. P. Spychalski, A. Błażyńska-Spychalska, J. Kobiela, Estimating case fatality rates of Covid-19. *Lancet Infect. Dis.* (2020).
3. J. H. Stock, "Data gaps and the policy response to the novel coronavirus" (Tech. Rep., National Bureau of Economic Research, 2020).
4. C. F. Manski, F. Molinari, "Estimating the Covid-19 infection rate: Anatomy of an inference problem" (Tech. Rep., National Bureau of Economic Research, 2020).
5. A. Atkeson, K. Kopecky, T. Zha, "Behavior and the transmission of Covid-19" in *Proceedings of AEA Papers* (2021).
6. I. Korolev, Identification and Estimation of the SEIRD Epidemic Model for COVID-19. Working Paper (2020).
7. D. Acemoglu, V. Chernozhukov, I. Werning, M. Whinston, "A multi-risk sir model with optimally targeted lockdown" (Tech. Rep., National Bureau of Economic Research, 2020).
8. A. Wald, Statistical decision functions (1950).
9. I. Gilboa, D. Schmeidler, Maxmin expected utility with non-unique prior. *J. Math. Econ.* **18**, 141-153 (1989).
10. L. Hansen, T. J. Sargent, Robust control and model uncertainty. *Am. Econ. Rev.* **91**, 60-66 (2001).
11. F. H. Knight, *Risk, Uncertainty and Profit*, 1921 (Library Economics Liberty, 1921).
12. K. J. Arrow, Alternative approaches to the theory of choice in risk-taking situations. *Econometrica* **19**, 404-437 (1951).
13. N. Bloom, The impact of uncertainty shocks. *Economics* **77**, 623-685 (2009).
14. S. R. Baker, N. Bloom, S. J. Davis, Measuring economic policy uncertainty. *Q. J. Econ.* **131**, 1593-1636 (2016).
15. C. Jones, T. Philippon, V. Venkateswaran, Optimal mitigation policies in a pandemic: Social distancing and working from home. *Rev. Financ. Stud.* **34**, 5188-5223 (2021).
16. E. W. Anderson, L. P. Hansen, T. J. Sargent, A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *J. Eur. Econ. Assoc.* **1**, 68-123 (2003).
17. F. Maccheroni, M. Marinacci, A. Rustichini, Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* **74**, 1447-1498 (2006).
18. L. P. Hansen, T. J. Sargent, Robustness and ambiguity in continuous time. *J. Econ. Theory* **146**, 1195-1223 (2011).
19. L. P. Hansen, J. Miao, Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 9163-9168 (2018).
20. I. Gilboa, S. Minardi, L. Samuelson, Theories and cases in decisions under uncertainty. *Games Econ. Behav.* **123**, 22-40 (2020).
21. D. Ellsberg, Risk, ambiguity, and the savage axioms. *Q. J. Econ.* **75**, 643-669 (1961).
22. F. J. Anscombe, R. J. Aumann, A definition of subjective probability. *Ann. Math. Stat.* **34**, 199-205 (1963).
23. L. J. Savage, *The Foundations of Statistics* (Courier Corporation, 1972).
24. P. J. Maenhout, Robust portfolio rules and asset pricing. *Rev. Financ. Stud.* **17**, 951-983 (2004).
25. L. Garlappi, R. Uppal, T. Wang, Portfolio selection with parameter and model uncertainty: A multi-prior approach. *Rev. Financ. Stud.* **20**, 41-81 (2007).
26. T. Cogley, R. Colacito, L. P. Hansen, T. J. Sargent, Robustness and U.S. monetary policy experimentation. *J. Money Credit Bank.* **40**, 1599-1623 (2008).
27. Y. Izhakian, D. Yermack, Risk, ambiguity, and the exercise of employee stock options. *J. Financ. Econ.* **124**, 65-85 (2017).
28. H. Ai, R. Bansal, H. Guo, A. Yaron, Identifying preference for early resolution from asset prices. Working Paper (2019).
29. J. Borovička, Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. *J. Polit. Econ.* **128**, 206-251 (2020).
30. D. M. Lemoine, C. P. Traeger, Tipping points and ambiguity in the economics of climate change. National Bureau of Economic Research, Working Paper 18230 (2012).
31. X. Li, B. Nezami Narajabad, T. P. Temzelides, Robust dynamic optimal taxation and environmental externalities (2014).
32. M. Barnett, W. Brock, L. P. Hansen, Pricing uncertainty induced by climate change. *Rev. Financ. Stud.* **33**, 1024-1066 (2020).
33. M. Barnett, W. Brock, L. P. Hansen, Climate change uncertainty spillover in the macroeconomy. *NBER Macroecon. Ann.* **36**, 387-395 (2021).
34. L. Berger *et al.*, Uncertainty and decision-making during a crisis: How to make policy decisions in the Covid-19 context? University of Chicago, Becker Friedman Institute for Economics Working Paper (2020).
35. A. B. Abel, S. Panageas, Optimal management of a pandemic in the short run and the long run. NBER Working Paper (2020).
36. G. Kaplan, B. Moll, G. Violante, "Pandemics according to hank" (Tech. Rep., Working Paper, 2020).
37. S. R. Baker, N. Bloom, S. J. Davis, S. J. Terry, "Covid-induced economic uncertainty" (Tech. Rep., National Bureau of Economic Research, 2020).
38. J. Kozlowski, L. Veldkamp, V. Venkateswaran, "Scarring body and mind: The long-term belief-scarring effects of Covid-19" (Tech. Rep., National Bureau of Economic Research, 2020).
39. M. S. Eichenbaum, S. Rebelo, M. Trabandt, "The macroeconomics of epidemics" (Tech. Rep., National Bureau of Economic Research, 2020).
40. F. E. Alvarez, D. Argente, F. Lippi, "A simple planning problem for Covid-19 lockdown" (Tech. Rep., National Bureau of Economic Research, 2020).
41. A. Atkeson, "How deadly is Covid-19? Understanding the difficulties with estimation of its fatality rate" (Tech. Rep., National Bureau of Economic Research, 2020).
42. C. Avery, W. Bossert, A. Clark, G. Ellison, S. F. Ellison, "Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists" (Tech. Rep., National Bureau of Economic Research, 2020).
43. J. Baek *et al.*, "The limits to learning a diffusion model" in *Proceedings of the 22nd ACM Conference on Economics and Computation* (2021), pp. 130-131.
44. N. Yamamoto, B. Jiang, H. Wang, Quantifying compliance with Covid-19 mitigation policies in the US: A mathematical modeling study. *Infect. Dis. Model.* **6**, 503-513 (2021).
45. J. N. Barrot, B. Grassi, J. Sauvagnat, Sectoral effects of social distancing. Working Paper (2020).
46. H. Wang *et al.*, Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China. *Cell Discov.* **6**, 1-8 (2020).
47. F. Arroyo-Marioli, F. Bullano, S. Kucinskas, C. Rondón-Moreno, Tracking R of Covid-19: A new real-time estimation using the Kalman filter. *PLoS One* **16**, e0244474 (2021).
48. H. Chernoff, A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Stat.* **23**, 493-507 (1952).
49. C. M. Newman, B. Stuck, Chernoff bounds for discriminating between two Markov processes. *Stochast.: Int. J. Probability Stochastic Processes* **2**, 139-153 (1979).
50. G. E. Box, Science and statistics. *J. Am. Stat. Assoc.* **71**, 791-799 (1976).