

NBER WORKING PAPER SERIES

EPIDEMIC RESPONSES UNDER UNCERTAINTY

Michael Barnett
Greg Buchak
Constantine Yannelis

Working Paper 27289
<http://www.nber.org/papers/w27289>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020, Revised March 2024

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. Constantine Yannelis is grateful to the Fama Miller Center for generous financial support. Finally, our paper benefited tremendously from the comments and suggestions from the editor and two anonymous referees. An online repository with the code, files, and results for our numerical analysis can be found at <https://github.com/mbarnet0/Covid>. 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). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Michael Barnett, Greg Buchak, and Constantine Yannelis. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Epidemic Responses Under Uncertainty
Michael Barnett, Greg Buchak, and Constantine Yannelis
NBER Working Paper No. 27289
May 2020, Revised March 2024
JEL No. E1,H0,I1

ABSTRACT

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.

Michael Barnett
W. P. Carey School of Business
Arizona State University
Department of Finance
300 E Lemon St
Tempe, AZ 85287
michael.d.barnett@asu.edu

Constantine Yannelis
Booth School of Business
University of Chicago
5807 S. Woodlawn Avenue
Chicago, IL 60637
and NBER
constantine.yannelis@chicagobooth.edu

Greg Buchak
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305b
buchak@stanford.edu

1 Introduction

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 (Chater, 2020). 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 (Spychalski, Błażyńska-Spychalska, and Kobiela, 2020). Similar difficulties¹ in estimating \mathcal{R}_0 led to estimates ranging from 1.5 to 12 (Korolev, 2020). To highlight the wide variation in these estimates, Figure 1 shows initial CFRs and \mathcal{R}_0 across many countries and US states.²


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. (2020), for example, state that “*We hasten to emphasize that there is considerable uncertainty about... the exact*

¹Many academic papers, e.g., Stock (2020), 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 and Molinari (2020) 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, Kopecky, and Zha (2021) provide a model explaining the range of estimates through behavioral responses.


²The insights of this model apply not only to new epidemic diseases, but also to new strains of existing diseases. While we focus for 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.

Figure 1: Estimated CFR Rates and \mathcal{R}_0

Panel A: Estimated CFR Rates by Country



Panel B: Estimated \mathcal{R}_0 for Countries and States



Panel A shows estimated CFRs for all countries with more than 1,000 cases and 100 deaths. Panel B shows estimated \mathcal{R}_0 across countries and states given different incubation periods (σ) and durations (γ). Source: European Centre for Disease Prevention and Control and [Korolev \(2020\)](#).

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 tradeoff 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 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, [Wald \(1950\)](#), [Gilboa and Schmeidler \(1989\)](#), and more recently [Hansen and Sargent \(2001\)](#) 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 towards 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 towards 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 towards 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.³ 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 [Knight \(1921\)](#) and [Arrow \(1951\)](#), 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.⁴ In our context, we introduce risk by allowing the disease to spread and kill non-deterministically. 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 his 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. An ambiguity-neutral planner with an equal-weighted prior across models⁵ adopts a relatively

³SIR 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. A recent literature in macroeconomics incorporates SIR models into macroeconomics models. [Stanford Earth System Sciences notes](#) provide an introduction to the standard epidemiological SIR model.

⁴A large body of literature refers to Brownian shocks and time variation in exposure to Brownian shocks as uncertainty, for example [Bloom \(2009\)](#) and [Baker, Bloom, and Davis \(2016\)](#). 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. [Jones et al. \(2021\)](#) provide a more complex model in which households respond to policy.

⁵We also explore three additional sets of prior beliefs over the models in the SI Appendix: (i) underestimating the pandemic, (ii) correctly estimating the pandemic, and (iii) overestimating the pandemic. These

modest quarantine, with the quarantine isolating roughly 50% of the population around 8 weeks into the pandemic and ending after roughly 15 weeks. In contrast, an ambiguity averse policymaker with the same prior immediately implements a quarantine policy of isolating 20% of the population, increasing it to 60% in 10 weeks, and maintaining some level of quarantine beyond 25 weeks. 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 push the planner to adopt more stringent quarantines, while uncertainty about the economic costs of quarantine push 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 a literature on ambiguity and robust control beginning with [Wald \(1950\)](#) and [Gilboa and Schmeidler \(1989\)](#), and continued by [Hansen and Sargent \(2001\)](#), [Anderson, Hansen, and Sargent \(2003\)](#), [Maccheroni, Marinacci, and Rustichini \(2006\)](#), [Hansen and Sargent \(2011\)](#), [Hansen and Miao \(2018\)](#), and [Gilboa, Minardi, and Samuelson \(2020\)](#).⁶ Recent work in finance and macroeconomics has emphasized the importance of different forms of ambiguity and uncertainty, for example [Maenhout \(2004\)](#), [Garlappi, Uppal, and Wang \(2007\)](#), [Cogley, Colacito, Hansen, and Sargent \(2008\)](#), [Bloom \(2009\)](#), [Baker, Bloom, and Davis \(2016\)](#), [Izhakian and Yermack \(2017\)](#), [Ai, Bansal, Guo, and Yaron \(2019\)](#), and [Borovička \(2020\)](#). Robust control methods have also been used to study the economic impacts of climate change, as in [Lemoine and Traeger \(2012\)](#), [Li, Nezami Nara-jabad, and Temzelides \(2014\)](#), [Barnett, Brock, and Hansen \(2020\)](#), and [Barnett et al. \(2021\)](#), where climate damages can have both permanent growth effects and transitory level effects.

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.

⁶There is an important and extensive body of theoretical study on uncertainty of various forms dating to [Knight \(1921\)](#), [Ellsberg \(1961\)](#), [Anscombe and Aumann \(1963\)](#), and [Savage \(1972\)](#).

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 [Abel and Panageas \(2020\)](#), [Kaplan, Moll, and Violante \(2020\)](#), [Jones, Philippon, and Venkateswaran \(2021\)](#), [Baker, Bloom, Davis, and Terry \(2020\)](#), [Kozlowski, Veldkamp, and Venkateswaran \(2020\)](#), [Eichenbaum, Rebelo, and Trabandt \(2020\)](#), and [Alvarez, Argente, and Lippi \(2020\)](#). These studies rely on calibrated parameters, which are often unknown. Parameter uncertainty is widely noted in this literature, and authors typically use a range of values. For example, [Acemoglu, Chernozhukov, Werning, and Whinston \(2020\)](#) note that: *“We stress that there is much uncertainty about many of the key parameters for COVID-19 ([Manski and Molinari, 2020](#)) and any optimal policy, whether uniform or not, will be highly sensitive to these parameters (e.g., [Atkeson \(2020\)](#), [Avery et al. \(2020\)](#), [Stock \(2020\)](#))). So our quantitative results are mainly illustrative and should be interpreted with caution.”* [Berger et al. \(2020\)](#) 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 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. Section 2 presents our model. Section 3 describes how to account for uncertainty. Section 4 presents simulation results and Section 5 concludes.

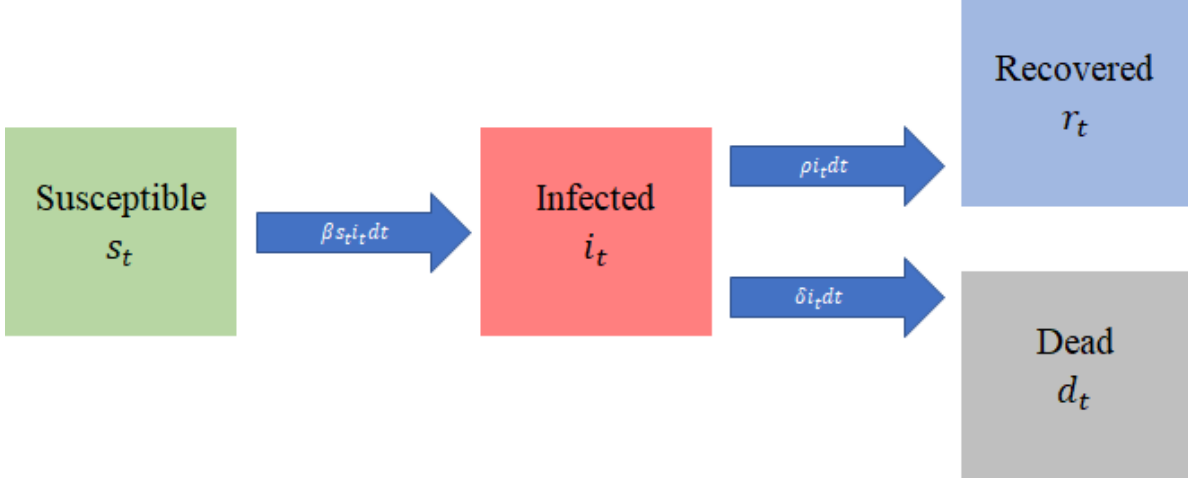
2 A Simple Economic Model of an Epidemic

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.

2.1 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

Figure 2: Expected Transition Rates Between States in the Augmented SIR Model.



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., $s_t = \frac{S_t}{N_t}$, $i_t = \frac{I_t}{N_t}$, $r_t = \frac{R_t}{N_t}$. 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.

Figure 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 increased infections have led to increased death rate due to limited resources for treatment resulting from increased hospitalizations. Similar to the frameworks used by [Alvarez et al. \(2020\)](#) and [Eichenbaum et al. \(2020\)](#), 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 ⁷. Two

⁷Formally, $W \doteq \{W_t : t \geq 0\}$ is a multi-dimensional Brownian motion where the corresponding filtration

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, σ_d , respectively. These shocks capture, for example, variability in exposure, co-morbidities, 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, the 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, r_t by way of Ito's lemma.

2.1.1 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 [Alvarez, Argente, and Lippi \(2020\)](#) 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.

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 .

2.2 Economic and Public Health Model

2.2.1 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 policymaker 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 by⁸

$$U_t = \kappa \log C_t$$

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 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 non-quarantined fraction of the available labor force.⁹

2.2.2 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

⁸We discuss the impact of including nonpecuniary losses for deaths from the pandemic as in [Alvarez et al. \(2020\)](#), [Abel and Panageas \(2020\)](#), and [Jones et al. \(2021\)](#) in the SI Appendix. Such costs, and uncertainty about these costs, should serve to enhance the results we find in our main analysis.

⁹An 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 qualitatively similar.

individuals who are not locked down may become less productive due to economic disruptions caused by lockdowns. To incorporate this additional channel, we assume 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:

$$A_t = \bar{A} \exp(z_t)$$

$$dz_t = (-\alpha q_t - z_t)dt + \sigma_z dW_z$$

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.

2.2.3 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 literature such as [Alvarez et al. \(2020\)](#) and [Abel and Panageas \(2020\)](#), 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 years. 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 the SI Appendix.

2.3 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.

3 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.¹⁰ The equilibrium solution has a Markovian structure such that the value

¹⁰We can omit the state variable r_t from our framework because the explicit adding up constraint of the SIR model means r_t is the residual after accounting for s_t, i_t , and N_t , and is therefore redundant.

function that characterizes the solution and the optimal quarantine policy are functions of the state variables s_t, i_t, N_t, 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” 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 [Hansen and Miao \(2018\)](#), and applied in the analysis of [Barnett, Brock, and Hansen \(2020\)](#). The result is a min-max 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 min-max problem optimization.

3.1 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

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

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 post-pandemic. 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

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

where we drop the θ notation for brevity. See the 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

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

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.

3.2 Optimal Policy with Uncertainty

To analyze the impact of model uncertainty, we implement a smooth ambiguity framework.¹¹ 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 [Hansen and Miao \(2018\)](#) for details about relative entropy in this setting. Using relative entropy means we are only considering relatively small distortions from the baseline model, but even small distortions can have sig-

¹¹The most common alternative to a smooth ambiguity approach is a robust preference approach as in [Anderson, Hansen, and Sargent \(2003\)](#). 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 the SI Appendix we provide an extension of the model where we apply the continuous-time robustness framework studied in [Hansen and Sargent \(2001\)](#), [Anderson, Hansen, and Sargent \(2003\)](#), [Maccheroni, Marinacci, and Rustichini \(2006\)](#), 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.

nificant 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

$$V(s_t, i_t, N_t, z_t) = \max_{q_t} \min_{\tilde{\pi}(\theta)} E_0[e^{-\kappa(T-t)} \hat{V}(N_T, z_T) + \int_0^T \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \{e^{-(\kappa+\lambda)t} \left(\kappa \log C(q_t) + \xi_a \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \right) \} dt]$$

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

$$\begin{aligned} (\kappa + \lambda)v(s_t, i_t) = & \max_{q_t} \min_{\tilde{\pi}(\theta)} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) \\ & + \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \left\{ -\frac{\alpha\kappa}{\kappa + 1} q_t + \delta_t [-i_t + (v_i i_t + v_s s_t) i_t] \right. \\ & + (v_i - v_s) \beta s_t i_t (1 - \zeta q_t)^2 - v_i i_t \gamma \left. \right\} - \frac{1}{2} i_t^2 \sigma_d^2 \\ & + \frac{1}{2} [v_{ss}(\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2 + v_{ii}(\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)] \end{aligned}$$

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

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 the optimally distorted probability weights are given by

$$\begin{aligned} \tilde{\pi}(s_t, i_t; \theta) \propto & \pi(\theta) \exp\left(-\frac{1}{\xi_a} \left\{ -\frac{\alpha(\theta)\kappa}{\kappa+1} q_t \right. \right. \\ & + \delta_t(\theta)(-i_t + v_s s_t i_t + v_i i_t^2) \\ & \left. \left. + \beta(\theta) s_t i_t (1 - \zeta(\theta) q_t)^2 (v_i - v_s) \right\} \right) \end{aligned}$$

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

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

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

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

Now that ambiguity is incorporated into the optimal policy choice, the planner tilts the probability weights towards 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. \(2021\)](#) 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.

4 Numerical Results

We now provide numerical results from simulations based on the theoretical solutions provided. Calibration of parameters is discussed first, then solutions for the models without uncertainty or “Outside the Model Uncertainty” sensitivity analysis, and then solutions for the model with uncertainty or “Inside the Model Uncertainty” sensitivity analysis.

4.1 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 non-pandemic version of the model ($A_0 \times L_0$) matches recent, pre-pandemic 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

the expected arrival time for a vaccine is 1.0 years, 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. \(2021\)](#) 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, Grassi, and Sauvagnat \(2020\)](#) estimate the GDP losses for numerous European countries based on 6-weeks 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 [Barrot et al. \(2020\)](#) estimates based on a fixed value of $q_t = 0.5$ over 6 weeks. 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 $\{0, 0.8\}$.

Table 1: Parameter Values

Parameter	Variable	Value
Subjective Discount Rate	κ	0.0025
Non-Pandemic 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	$\{\infty, 0.0032\}$

For the pandemic model parameters, we use values from various studies (including [Korolev \(2020\)](#), [Atkeson \(2020\)](#), [Wang, Wang, Dong, Chang, Xu, Yu, Zhang, Tsamlag, Shang, Huang, et al. \(2020\)](#), [Abel and Panageas \(2020\)](#), 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 . The value of γ is held fixed at $\gamma = \frac{30}{18}$, or an expected duration of infection of 2.5 weeks. 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 add 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 [Arroyo-Marioli et al. \(2021\)](#), 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 days of the COVID 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 value¹². 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.¹³ For σ_z , we use a value that matches [Barnett et al. \(2020\)](#), who calculate time series standard deviations 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 .

¹²We 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.

¹³While 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 .

4.2 Outside-the-Model Uncertainty: Sensitivity Analysis

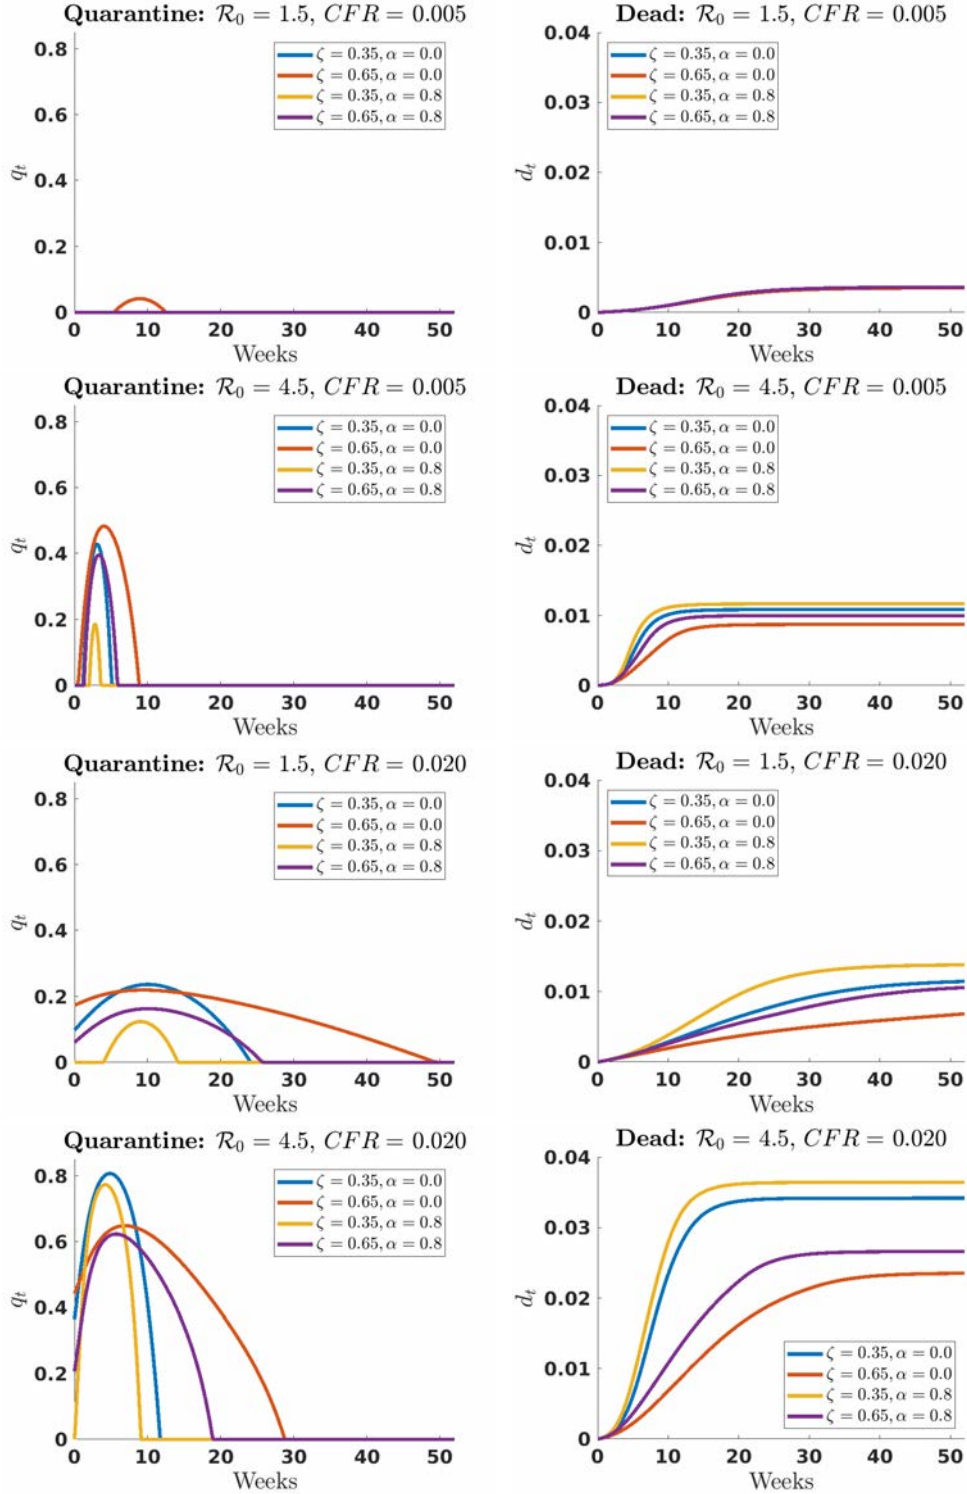
We first provide simulated outcomes of the model based on different pandemic models without the planner accounting for uncertainty in their optimal decision. This corresponds to what is typically termed as a sensitivity analysis and illustrates the wide range of optimal responses that depend on the underlying model parameters. Figure 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 $\alpha \in \{0, 0.8\}$.

Figure 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 leads 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 zero weeks to approaching one year. The fraction of dead in the population resulting from these policies varies by an order of magnitude, running between nearly 0.0025% to 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 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.

4.3 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 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

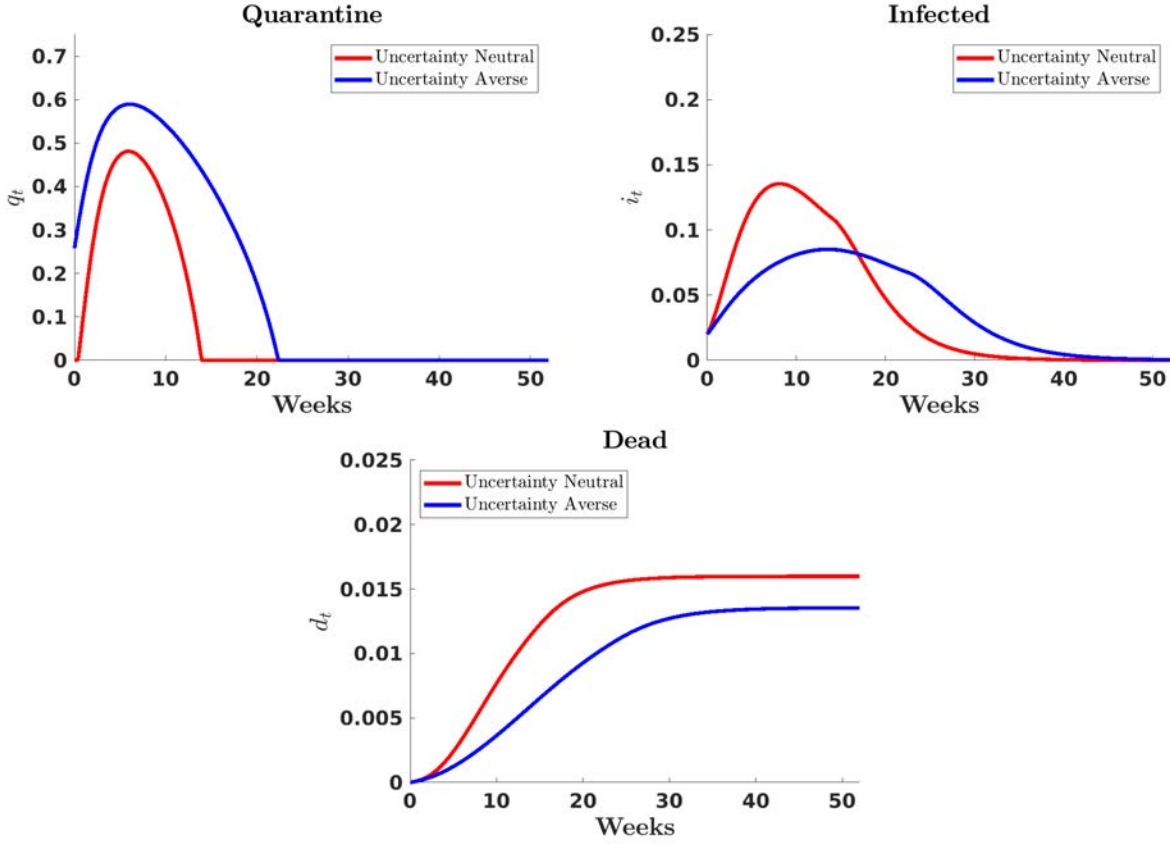
Figure 3: Outside the Model Uncertainty



Notes: These figures show the range of outcomes and policy responses across 16 potential models of the pandemic that vary by $\mathcal{R}_0, CFR, \alpha$, 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 $\zeta \in \{0.35, 0.65\}$.

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 the 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.

Figure 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 the uncertainty averse model response (blue).

Figure 4 shows the infected, dead, and quarantine outcomes for the uncertainty neutral and uncertainty averse cases of the equal-weighted prior scenario. Figure 5 shows how the uncertainty averse planner's prior probabilities over each model are adjusted and distorted over time. Figure 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 Figure 4, we see that quarantine policy starts at 0, then quickly moves up to around 50% near week 8, and then quickly drops back to 0 by week 15. The underlying pandemic peaks at around 15% infected near week

8 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 weeks). The resulting impact on the pandemic is that infections never go above 10%, and deaths are reduced by between a quarter to a half a percentage point.

Figure 5: Uncertainty Aversion Distorted Model Probabilities

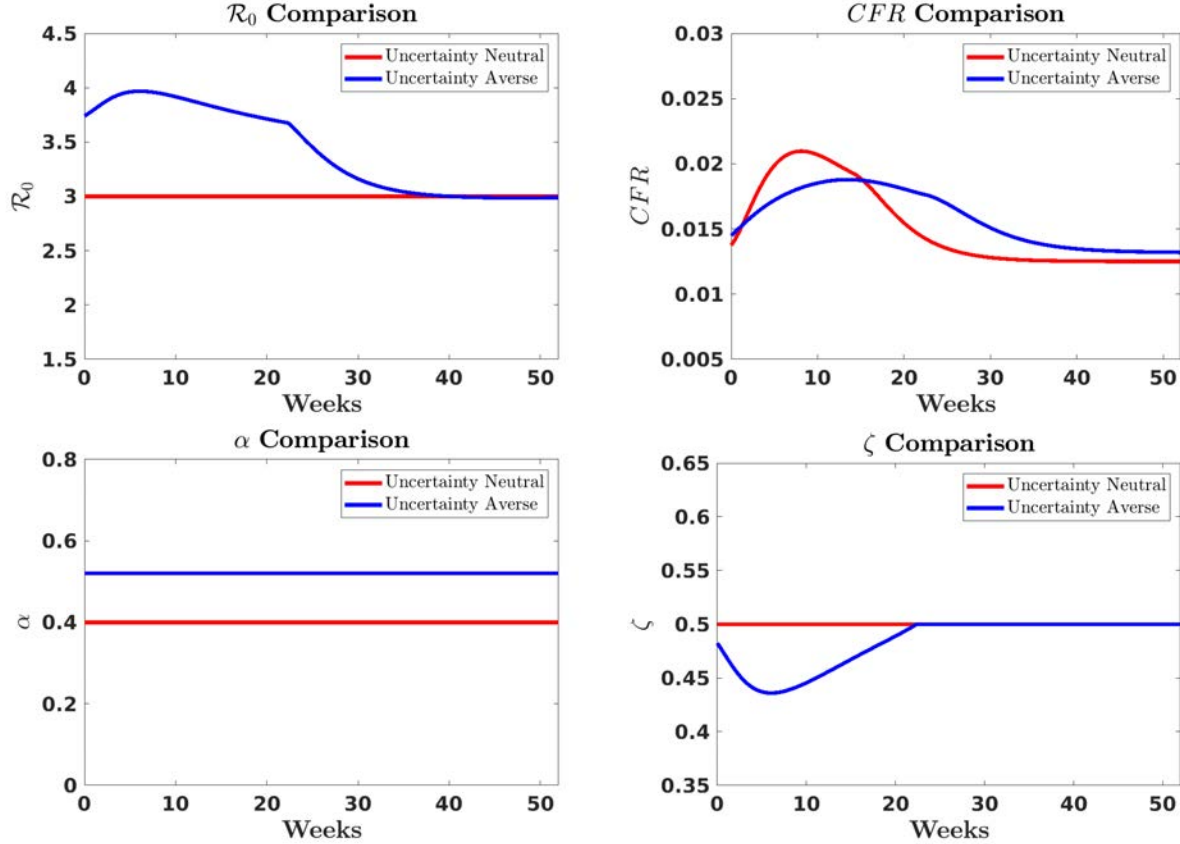


Notes: This figure shows the distorted probability weights for the uncertainty averse planner. The dashed lines are the probability weights on the most severe epidemiological scenario ($CFR = 0.02$, $\mathcal{R}_0 = 4.5$) with the least effective pandemic mitigation scenario ($\zeta = 0.35$). The red 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.

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 Figure 5 and the distorted model parameters of the uncertainty averse planner in Figure 6 highlights why this is the case. The two models in Figure 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 (the yellow and red dashed lines). And while the model of these two that also assumes the highest economic costs of quarantine (the 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 (the 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 (the red shaded region) remain relatively low throughout, with some dropping to allow for a shift in probability to the worst-case models.

Figure 6: Distorted Model Parameters: $\mathcal{R}_0, CFR, \alpha, \zeta$



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.

Figure 6 shows the distorted model parameters implied by the planner's aversion to model ambiguity (the blues lines) compared with the true model parameters (the 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 week 25. 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


tions 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 Figure 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 Figure 7, for this. These probabilities, now commonly used in the literature on model uncertainty, are based on the model discrimination methods proposed by Chernoff (1952) and Newman and Stuck (1979). 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. Under a standard heuristic, a value below 10% or 5% would indicate that a policymaker should reject the distorted model in light of data generated by the baseline model.¹⁴ Figure 7 shows clear dynamic shifts in the detection error probabilities for the uncertainty averse, with the probability starting just above 10% and then converging towards 50% after 52 weeks as infections converge towards zero. At that point, when the pandemic has largely run its course and the distorted weights are converging back towards 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,

¹⁴Following Anderson et al. (2003), the detection error probabilities are calculated using a log-likelihood ratio based model selection criteria that compares model outcomes for repeated simulations using the baseline and worst-case models.

Figure 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).

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.

5 Concluding Remarks

In this paper, we embrace the admonition of [Box \(1976\)](#), 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 only had small outbreaks, there is often significant uncertainty about key parameters which determine the overall consequences of an epidemic.

The uncertainty averse planner tilts the prior towards 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 criteria. 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.

References

- Abel, A. B. and S. Panageas (2020). Optimal management of a pandemic in the short run and the long run. *NBER Working Paper* (w27742).
- Acemoglu, D., V. Chernozhukov, I. Werning, and M. Whinston (2020). A multi-risk sir model with optimally targeted lockdown. Technical report, National Bureau of Economic Research.
- Ai, H., R. Bansal, H. Guo, and A. Yaron (2019). Identifying preference for early resolution from asset prices. *Working Paper*.
- Alvarez, F. E., D. Argente, and F. Lippi (2020). A simple planning problem for covid-19 lockdown. Technical report, National Bureau of Economic Research.
- Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association* 1(1), 68–123.
- Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *Annals of Mathematical Statistics* 34(1), 199–205.
- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica: Journal of the Econometric Society*, 404–437.
- Arroyo-Marioli, F., F. Bullano, S. Kucinskas, and C. Rondón-Moreno (2021). Tracking r of covid-19: A new real-time estimation using the kalman filter. *PloS one* 16(1), e0244474.
- Atkeson, A. (2020). How deadly is covid-19? understanding the difficulties with estimation of its fatality rate. Technical report, National Bureau of Economic Research.
- Atkeson, A., K. Kopecky, and T. Zha (2021). Behavior and the transmission of covid-19. *AEA Papers and Proceedings*.
- Avery, C., W. Bossert, A. Clark, G. Ellison, and S. F. Ellison (2020). Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists. Technical report, National Bureau of Economic Research.
- Baek, J., V. F. Farias, A. Georgescu, R. Levi, T. Peng, D. Sinha, J. Wilde, and A. Zheng (2021). The limits to learning a diffusion model. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 130–131.
- Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics* 131(4), 1593–1636.
- Baker, S. R., N. Bloom, S. J. Davis, and S. J. Terry (2020). Covid-induced economic uncertainty. Technical report, National Bureau of Economic Research.
- Barnett, M., W. Brock, and L. P. Hansen (2020). Pricing uncertainty induced by climate change. *The Review of Financial Studies* 33(3), 1024–1066.
- Barnett, M., W. Brock, and L. P. Hansen (2021). Climate Change Uncertainty Spillover in the Macroeconomy. *NBER Macroeconomics Annual* 36(1), 387–395.
- Barrot, J.-N., B. Grassi, and J. Sauvagnat (2020). Sectoral effects of social distancing. *Working Paper*.
- Berger, L., N. Berger, V. Bosetti, I. Gilboa, L. P. Hansen, C. Jarvis, M. Marinacci, and R. Smith (2020). 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* (2020-95).

- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica* 77(3), 623–685.
- Borovička, J. (2020). Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. *Journal of Political Economy* 128(1), 206–251.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association* 71(356), 791–799.
- Cagetti, M., L. P. Hansen, T. Sargent, and N. Williams (2002). Robustness and pricing with uncertain growth. *The Review of Financial Studies* 15(2), 363–404.
- Chater, N. (2020). Facing up to the uncertainties of covid-19. *Nature Human Behaviour*, 1–1.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics* 23, 493–507.
- Cogley, T., R. Colacito, L. P. Hansen, and T. J. Sargent (2008). Robustness and us monetary policy experimentation. *Journal of Money, Credit and Banking* 40(8), 1599–1623.
- Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020). The macroeconomics of epidemics. Technical report, National Bureau of Economic Research.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 643–669.
- Garlappi, L., R. Uppal, and T. Wang (2007). Portfolio selection with parameter and model uncertainty: A multi-prior approach. *The Review of Financial Studies* 20(1), 41–81.
- Gilboa, I., S. Minardi, and L. Samuelson (2020). Theories and cases in decisions under uncertainty. *Games and Economic Behavior*.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Hansen, L. and T. J. Sargent (2001). Robust control and model uncertainty. *American Economic Review* 91(2), 60–66.
- Hansen, L. P. and J. Miao (2018). Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proceedings of the National Academy of Sciences* 115(37), 9163–9168.
- Hansen, L. P. and T. J. Sargent (2010). Fragile beliefs and the price of uncertainty. *Quantitative Economics* 1(1), 129–162.
- Hansen, L. P. and T. J. Sargent (2011). Robustness and Ambiguity in Continuous Time. *Journal of Economic Theory* 146(3), 1195–1223.
- Hansen, L. P., T. J. Sargent, G. Turmuhambetova, and N. Williams (2006). Robust Control and Model Misspecification. *Journal of Economic Theory* 128(1), 45–90.
- Hansen, L. P., T. J. Sargent, and N. E. Wang (2002). Robust permanent income and pricing with filtering. *Macroeconomic dynamics* 6(1), 40–84.
- Izhakian, Y. (2020). A theoretical foundation of ambiguity measurement. *Journal of Economic Theory* 187, 105001.
- Izhakian, Y. and D. Yermack (2017). Risk, ambiguity, and the exercise of employee stock options. *Journal of Financial Economics* 124(1), 65–85.
- Jones, C., T. Philippon, and V. Venkateswaran (2021). Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies* 34(11), 5188–5223.

- Kaplan, G., B. Moll, and G. Violante (2020). Pandemics according to hank. Technical report, Working Paper.
- Knight, F. H. (1921). Risk, uncertainty and profit, 1921. *Library of Economics and Liberty*.
- Korolev, I. (2020). Identification and Estimation of the SEIRD Epidemic Model for COVID-19. *Working Paper*.
- Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020). Scarring body and mind: the long-term belief-scarring effects of covid-19. Technical report, National Bureau of Economic Research.
- Kushner, H. J. and P. Dupuis (2001). *Numerical Methods for Stochastic Control Problems in Continuous Time*, Volume 24. Springer Science & Business Media.
- Lemoine, D. M. and C. P. Traeger (2012, July). Tipping Points and Ambiguity in the Economics of Climate Change. Working Paper 18230, National Bureau of Economic Research.
- Li, X., B. Nezami Narajabad, and T. P. Temzelides (2014). Robust dynamic optimal taxation and environmental externalities.
- Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6), 1447–1498.
- Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of Financial Studies* 17(4), 951–983.
- Manski, C. F. and F. Molinari (2020). Estimating the covid-19 infection rate: Anatomy of an inference problem. Technical report, National Bureau of Economic Research.
- Newman, C. M. and B. Stuck (1979). Chernoff bounds for discriminating between two markov processes. *Stochastics: An International Journal of Probability and Stochastic Processes* 2(1-4), 139–153.
- Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.
- Spychalski, P., A. Błażyńska-Spychalska, and J. Kobiela (2020). Estimating case fatality rates of covid-19. *The Lancet Infectious Diseases*.
- Stock, J. H. (2020). Data gaps and the policy response to the novel coronavirus. Technical report, National Bureau of Economic Research.
- Wald, A. (1950). Statistical decision functions.
- 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.
- Yamamoto, N., B. Jiang, and H. Wang (2021). Quantifying compliance with covid-19 mitigation policies in the us: A mathematical modeling study. *Infectious Disease Modelling* 6, 503–513.

Appendix A SIR Model Dynamics and Volatility

A traditional SIR model is defined by state variables for the number of susceptible individuals S_t , the number of infected individuals I_t , and the number of recovered individuals R_t . In addition, the total population is tracked by the state variable N_t . The evolution of these state variables is typically determined by the deterministic differential equations

$$\begin{aligned} dS_t &= \psi_t N_t dt - \omega_t S_t dt - \beta_t S_t \frac{I_t}{N_t} dt \\ dI_t &= \beta_t S_t \frac{I_t}{N_t} dt - \omega_t I_t dt - \gamma_t I_t dt \\ dR_t &= -\omega_t R_t dt + \rho_t I_t dt \\ dN_t &= \psi_t N_t dt - \omega_t (S_t + I_t + R_t) dt - \delta_t I_t dt \end{aligned}$$

where β_t is the pandemic infection rate, ρ_t is the pandemic recovery rate, δ_t is the pandemic death rate, $\gamma_t = \rho_t + \delta_t$ is the expected duration of infection, ψ_t is the birth rate, and ω_t is the non-pandemic related death rate. In addition there is an adding up constraint such that $N_t = S_t + I_t + R_t$.

We then define new state variables, $s_t = \frac{S_t}{N_t}$, $i_t = \frac{I_t}{N_t}$, and $r_t = \frac{R_t}{N_t}$ as the susceptible, infected, and recovered fractions of the total population. Applying Ito's lemma to these new state variables, given the original level state variables dynamics, we find the evolution of the fraction of the total population state variables as

$$\begin{aligned} ds_t &= \psi_t dt - \beta_t s_t i_t dt - s_t(\psi_t - i_t \delta_t) dt \\ di_t &= \beta_t s_t i_t dt - (\rho_t + \delta_t) i_t dt - i_t(\psi_t - i_t \delta_t) dt \\ dr_t &= \rho_t i_t dt - r_t(\psi_t - i_t \delta_t) dt \end{aligned}$$

Finally to arrive at the evolution equations we use in our analysis we make two assumptions. First, we assume $\omega_t = \psi_t = 0$, allowing us to focus on the dynamics associated with the pandemic. Second, we specify functional forms and add volatility in the form of parameter perturbations so that we replace $\beta_t dt$ by $\beta dt + \sigma_i dW_i$ and $\delta_t dt$ by $\delta_t dt + \sigma_d dW_d$. This results in the state variable evolution equations given in the beginning of the main text. Furthermore, adding in quarantine impacts as discussed in the paper gives us the final version of the state variable evolution equations used in our analysis.

Appendix B Detection Error Probabilities

To quantify the magnitude of uncertainty allowed for in our analysis based on the choice of ξ_a , we use detection error probabilities as proposed by Hansen et al. (2002) and Anderson et al. (2003). These probabilities, now commonly used in the literature on model uncertainty, are based on the model discrimination bounds proposed by Chernoff (1952) and Newman and Stuck (1979). The bound quantifies the probability of making a type I or type II error when choosing between two models. In our setting, the two models are the baseline model where state dynamics depend on the prior probabilities and the worst-case model where state dynamics depend on the distorted posterior.

We construct this probability bound in the smooth ambiguity setting by recasting the density distortion from aversion to model ambiguity as a drift distortion. In the setting of misspecification concerns and robust preferences, this drift distortion arises from the decision theoretic structure because the Brownian motion dW_t is replaced by $h_t dt + d\hat{W}_t$ in the model of state dynamics. The distortions are disguised by the Brownian shocks dW_t , making them difficult to distinguish, and are without parametric form. The drift distortion in our framework depends on the structured parameter uncertainty, as we show next.

Consider a general state vector \mathbf{x}_t where, under a baseline prior, the state evolves as

$$d\mathbf{x}_t = \boldsymbol{\mu} dt + \boldsymbol{\sigma} dW_t$$

and under the distorted prior, the evolution is given by

$$d\mathbf{x}_t = \tilde{\boldsymbol{\mu}} dt + \boldsymbol{\sigma} dW_t$$

Note that $\boldsymbol{\sigma}$ is the matrix of volatilities for each state and shock and $\boldsymbol{\mu}$ and $\tilde{\boldsymbol{\mu}}$ are the vector of drifts for each state variable weighted by the prior weights $\pi(\theta)$ and the uncertainty adjusted weights $\tilde{\pi}(\theta)$ across each model θ , respectively. Following Barnett et al. (2020), the drift distortion implied by the structured ambiguity across model parameters is given by

$$\mathbf{h}_t = (\boldsymbol{\sigma}' \boldsymbol{\sigma})^{-1} \boldsymbol{\sigma}' (\tilde{\boldsymbol{\mu}} - \boldsymbol{\mu})$$

Given the drift distortions, we solve for the detection error probability through simulation methods as in Anderson et al. (2003), Hansen et al. (2002), and Hansen and Sargent (2010). This provides us with a state-dependent, conditional model discrimination measure to determine whether our choice of ξ_a is reasonable in that the detection error probability would be small enough that we would be confident in a statistically significant sense of which

model is the true model (i.e., $< 10\%$). To calculate the detection error probabilities, we use discrete-time approximations of the state variables of interest. Define $x_t \equiv [s_t, i_t, N_t, z_t]$ and $\tilde{x}_t \equiv [\tilde{s}_t, \tilde{i}_t, \tilde{N}_t, \tilde{z}_t]$, where x_t evolves under the baseline prior and \tilde{x}_t evolves under the distorted distribution as follows:

$$\begin{aligned}\mathbf{x}_{t+1} &= \boldsymbol{\mu}dt + \boldsymbol{\sigma}\sqrt{dt}W_{t+1} \\ \tilde{\mathbf{x}}_{t+1} &= \tilde{\boldsymbol{\mu}}dt + \boldsymbol{\sigma}\sqrt{dt}\tilde{W}_{t+1}\end{aligned}$$

W_{t+1} and \tilde{W}_{t+1} are matrices of shocks that follow a Standard Normal distribution. We then implement the following algorithm to construct the detection error probabilities:

1. Simulate many pathways for \mathbf{x}_{t+1} for a predetermined number of periods T .
2. Simulate many pathways for $\tilde{\mathbf{x}}_{t+1}$ for a predetermined number of periods T .
3. Calculate the log-likelihood ratios r_A and r_B as follows:
 - (a) Solve for $\check{W}_{t+1} = W_{t+1} - \check{h}_t$ where \check{h}_t is derived as shown above based on the pathways for x_{t+1} .
 - (b) Solve for $\hat{W}_{t+1} = \tilde{W}_{t+1} + \hat{h}_t$ where \hat{h}_t is derived as shown above based on the pathways for \tilde{x}_{t+1} .
 - (c) Calculate $r_A = \frac{1}{2T} \sum \{(W_{t+1} - \check{h}_t)'(W_{t+1} - \check{h}_t) - W_{t+1}'W_{t+1}\}$
 - (d) Calculate $r_B = \frac{1}{2T} \sum \{(\tilde{W}_{t+1} + \hat{h}_t)'(\tilde{W}_{t+1} + \hat{h}_t) - \tilde{W}_{t+1}'\tilde{W}_{t+1}\}$
4. Use the log-likelihood ratios to calculate $p_i = \text{freq}(r_i \leq 0), i = A, B$.
5. Derive the detection error probability as $p(\theta) = \frac{1}{2}(p_A + p_B)$.

Note that we set $T = 12$ months as mentioned previously, and we simulate 100,000 pathways for \mathbf{x}_{t+1} and $\tilde{\mathbf{x}}_{t+1}$ at each point in time of our simulated model solution results. Therefore, the detection error probability calculated at each point in time of our simulated model results is based on information from 100,000 independent and uniquely calculated pathways with equal sample length. The value of $1/2$ in step 5 comes from choosing an equal-weighted prior on each type of error p_A and p_B , where p_A represents the probability of choosing the distorted model as the model that generated the sample path when it is actually the baseline model and p_B represents the probability of choosing the baseline model as the model that generated the sample path when it is actually the distorted model. Thus, a detection error probability of 0.5 means that there is essentially no way to statistically discriminate between the distorted model and the baseline model, while a detection error

probability close to 0.0 means that one can statistically discriminate between the distorted model and the baseline model with near certainty.

Appendix C Model Extensions

Preferences with Nonpecuniary Losses from Deaths

We can extend our model by adding an additional cost and uncertainty component from the pandemic in the form of nonpecuniary losses from deaths due to the pandemic. To do this, we now assume the representative household has flow utility that depends on consumption C_t , a subjective discount rate κ , and nonpecuniary losses for deaths from the pandemic

$$U_t = \kappa \log C_t - x_t$$

The nonpecuniary losses for deaths from the pandemic x_t account for losses beyond the economic costs of a reduced labor force, and therefore reduced final output production, from quarantine measures and death. Because these costs are determined by deaths from the pandemic, there is a mapping between infections and nonpecuniary costs of the form $x_t = \chi(i_t)$, where we assume a functional form similar to [Alvarez et al. \(2020\)](#), [Abel and Panageas \(2020\)](#), and [Jones et al. \(2021\)](#):

$$\chi(\iota) = \varphi \delta_t \iota$$

where $\delta_t = \delta + \delta_+ \iota$ as noted in the main text. The value of φ represents the valuation the planner places on the deaths from the pandemic and $\delta_t \iota_t$ is the fraction of the population who die in any given period because of the pandemic. Applying Ito's lemma we derive the evolution of nonpecuniary costs as

$$dx_t = \chi_i \mu_i(s_t, i_t; q_t, \theta) dt + \frac{1}{2} \chi_{ii} |\sigma_i(s_t, i_t)|^2 dt + \chi_i \sigma_i(s_t, i_t) dW_t$$

where χ_i and χ_{ii} are the first and second derivatives, respectively of $\chi(i_t)$ and $\mu_i(s_t, i_t; q_t, \theta)$ and $\sigma_i(s_t, i_t)$ are the drift and volatility, respectively, of the state variable for infections i_t . Note optimal quarantine policy influences the drift of i_t which directly enters the evolution of x_t . Therefore, the optimal choice for q_t will include an explicit adjustment for these non-pecuniary costs. Also, we note that uncertainty will enter through the existing uncertainty related to \mathcal{R}_0 and the CFR , which enter into the drift for x_t . As such, this extension will serve to amplify the existing channels related to uncertainty and its implications for the

choice of optimal quarantine policy.

Uncertainty Through Robustness

Though our main analysis used the smooth ambiguity framework, where the social planner optimally chose probability weights to place on competing parameterizations of the model, an alternate approach to the problem is through applying the robust preferences methodology established in the economics literature¹⁵. Accounting for uncertainty in this way allows the social planner to make optimal mitigation policy choices while acknowledging that a given baseline model may be misspecified. As with smooth ambiguity, the mathematical tractability of the robust preferences decision problem allows us to characterize the implications of uncertainty for optimal policy decisions with clear intuition. We briefly outline here how we incorporate robust preferences to account for model uncertainty, and direct readers to the aforementioned references for complete mathematical details.

We define the approximating or baseline model using the evolution equations of the state variables as previously given:

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

As was the case in the smooth ambiguity setting, we assume the baseline model is the result of historical data or previous information about coronavirus pandemics and acts as a best-guess at what the true COVID-19 pandemic model is for policymakers. However, we allow the social planner in our model to consider the likelihood that this model is misspecified, or that there are possibly other models which are the true model for the COVID-19 pandemic.

Possible alternative models are represented by a drift distortion that is added to the approximating model by changing the Brownian motion W_t to $\hat{W}_t + \int_0^t h_s ds$ where h_s and \hat{W}_t are processes adapted to the filtration generated by the Brownian motion W_t . Therefore, alternative models under consideration by the social planner are of the form

$$ds_t = -\beta s_t i_t (1 - \zeta q_t)^2 dt + s_t i_t \delta_t dt - s_t i_t \sigma_i (h_t + d\hat{W}_i) + s_t i_t \sigma_d (h_t + d\hat{W}_d)$$

¹⁵Detailed explanations of robust preference problems and axiomatic treatment of such formulations using penalization methods are given by [Cagetti, Hansen, Sargent, and Williams \(2002\)](#), [Anderson, Hansen, and Sargent \(2003\)](#), [Hansen, Sargent, Turmuhambetova, and Williams \(2006\)](#), [Maccheroni, Marinacci, and Rustichini \(2006\)](#), [Izhakian \(2020\)](#) and [Hansen and Sargent \(2011\)](#).

$$\begin{aligned}
di_t &= \beta s_t i_t (1 - \zeta q_t)^2 dt - \gamma i_t dt + i_t^2 \delta_t dt + \sigma_i s_t i_t (h_t + d\hat{W}_i) - i_t \sigma_d (h_t + d\hat{W}_d) + i_t^2 \sigma_d (h_t + d\hat{W}_d) \\
dz_t &= -\alpha q_t dt - z_t dt + \sigma_z (h_t + d\hat{W}_z) \\
\frac{dN_t}{N_t} &= -i_t \delta_t dt - i_t \sigma_d (h_t + d\hat{W}_d)
\end{aligned}$$

In this form, the alternative models are disguised by the Brownian motion and so are hard to detect statistically using past data. In addition, the alternative models are given without direct parametric form, which allows for a larger class of alternative models under consideration by the planner. The h_t in the model will be optimally determined and state dependent, and so the magnitude of the parameter misspecification considered by the social planner when making optimal policy decisions will depend on the current state of the pandemic and evolve dynamically.

For the uncertainty analysis to be reasonable, we restrict the set of alternative models considered by the social planner to those that are difficult to distinguish from the baseline model using statistical methods and past data. A penalization term based on the conditional relative entropy measure of model distance is used to accomplish this. The parameter ξ_m is chosen to determine the magnitude of this penalization. We have defined relative entropy previously, and note that [Hansen, Sargent, Turmuhambetova, and Williams \(2006\)](#) provides complete details about relative entropy use in a robust preferences setting. Again, relative entropy means we are only considering relatively small, though potentially significant, distortions from the baseline model.

The time derivative of relative entropy or contribution of the current worst-case model $h_t dt$ to relative entropy is given by $\frac{1}{2}|h_t|^2$. This term is added to the flow utility or preferences of the household to account for model uncertainty. As was the case in the smooth ambiguity setting, optimal decisions will be determined by considering alternative worst-case models as a device to generate optimal policies that are robust to alternative models, and not as some type of distorted beliefs setting. The household maximization problem is replaced with a max-min set-up, where the minimization is made over possible model distortions h_t^* which are constrained by ξ_m . This allows the planner to determine the relevant worst-case model for given states of the world to help inform their optimal policy decisions.

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

$$V(s_t, i_t, N_t, z_t) = \max_{q_t} \min_{h_t} E_0 \left[\int_0^T e^{-(\kappa+\lambda)t} \{ \kappa \log(C_t(q_t)) + \frac{\xi_m}{2} |h_t|^2 \} dt + e^{-\kappa(T-t)} \tilde{V}(N_T, z_T) \right]$$

subject to market clearing and labor supply constraints.

As before, the social planner's solution is still 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

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

The first-order conditions for the optimal model distortions give us

$$|h_t|^2 = \frac{1}{\xi_m^2} [(v_i - v_s)^2 (s_t i_t \sigma_i)^2 + (s_t v_s - (1 - i_t) v_i - 1)^2 (i_t \sigma_d)^2 + \left(\frac{\kappa}{\kappa + 1} \right)^2 \sigma_z^2]$$

Plugging back in to the HJB equation, we are left with the following problem

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

The optimal choice of mitigation q_t is of the same functional form as in the smooth ambiguity analysis. Key differences to the social planner problem and HJB equation show up through the adjustments to the flow utility as a result of the penalization term accounting for model uncertainty concerns. The implications of model uncertainty for optimal mitigation policy and social welfare in the face of a pandemic are not only the direct adjustments to the key equations of interest, but also how these adjustments feed through the model solution and alter the value function V and the marginal values of changes to the susceptible, infected, total population, and productivity represented by V_s, V_i, V_N, V_z . Though we do not report results from this approach, they can easily be solved for numerically.

Uncertainty Through Alternative Smooth Ambiguity Formulation

Here we incorporate uncertainty using the decision theoretic framework developed in [Hansen and Sargent \(2011\)](#). Given an optimal policy $q_t(\theta)$ and value function $v(s_t, i_t; \theta)$ for each conditional model, we first specify a prior distribution to 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 reasonable. This works by restricting the set of alternative model weights considered by the social planner to those that are difficult to distinguish from the prior model distribution using statistical methods. The parameter ξ_a is chosen to determine the magnitude of this penalization. Large values of ξ_a imply low aversion to ambiguity, while small values of ξ_a imply strong aversion to ambiguity. Relative entropy is 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 two models.¹⁶

This new, second-stage problem for the planner is a minimization problem, where the minimization is made over possible distorted probability weights $\tilde{\pi}(\theta)$ which are constrained by ξ_a based on the solutions to the θ conditional value function solutions found previously. This allows the planner to determine the relevant worst-case model for given states of the world to help inform their optimal policy decisions.¹⁷ Though optimal decisions will be determined by considering alternative worst-case models, this setting should not be interpreted as a distorted beliefs model. The worst-case model is used as a device to produce solutions

¹⁶See [Hansen and Sargent \(2011\)](#) for details about relative entropy in this setting. Using relative entropy means we are only considering relatively small distortions from the baseline model, but even small distortions can have significant impacts on optimal policy. In particular, we apply relative entropy penalization directly to the set of conditional value functions.

¹⁷One can view the optimal policies choices as being made by a sequence of policymakers at each point in time, which assumes limited commitment in our framework. Issues of dynamic consistency and limited commitment are relevant for a broad class of optimal control problems solved under uncertainty. The limited commitment assumed by this interpretation of our model creates a possible tension in terms of whether the planner's optimal choices are consistent with time zero choices made. However, given that our discount factor is close to zero because of the weekly time-scale used for our analysis, the intertemporal differences in social valuations will be quite small and so we suspect that the impact of limited commitment will be quantitatively small.

that are robust to alternative models. The second-stage minimization problem is given by the solution to the following problem

$$\begin{aligned} \tilde{V}_t = \min_{\tilde{\pi}(\theta)} \sum_{\theta \in \Theta} \tilde{\pi}(\theta) (V(\theta) + \xi_a [\log(\tilde{\pi}(\theta)) - \log(\pi(\theta))]) \\ \text{subject to } \sum_{\theta \in \Theta} \pi(\theta) = \sum_{\theta \in \Theta} \tilde{\pi}(\theta) = 1 \end{aligned}$$

Taking the first order condition for this problem, and imposing $\sum \tilde{\pi}(\theta) = 1$, we find the optimally distorted probability weights are given by

$$\tilde{\pi}(\theta) = \pi(\theta) \frac{\exp(-\frac{1}{\xi_a} V(\theta))}{\sum \pi(\theta) \exp(-\frac{1}{\xi_a} V(\theta))}$$

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

From the distorted probability weights, we see that while the prior probability weights anchor the outcomes to a baseline expectation of the true model, smooth ambiguity leads to an exponential tilting towards those θ conditional models that lead to the most negative lifetime expected utility implications. In order to determine the ambiguity robust policy for the social planner, we weight the θ conditional optimal mitigation policies $q_t(\theta)$ using the distorted probability weights. The magnitude of the weight given to each θ conditional model informs the planner on how to use the θ conditional mitigation policies to determine an ambiguity robust optimal policy. This same re-weighting using the distorted probability weights provides us with the distorted parameters which the social planner uses to make optimal policy decisions in this setting. The adjusted mitigation policy and distorted parameters are therefore given by

$$\tilde{q}_t = \sum_{\theta \in \Theta} \tilde{\pi}(\theta) q_t(\theta), \quad \tilde{\beta}_t = \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \beta(\theta), \quad \tilde{\delta}_t = \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \delta_t(\theta) \quad \tilde{\zeta}_t = \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \zeta(\theta) \quad \tilde{\alpha}_t = \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \alpha(\theta)$$

As the planner tilts their value function and probability weights towards certain models, this leads to the implied distorted model parameters which are adjusted by worst-case outcomes which the planner uses as a lens to view and respond in a robustly optimal way in the face of uncertainty.

Appendix D Numerical Method

We solve numerically for a solution to the PDE representing the HJB equation for our social planner's problem. The numerical method we use to derive our solution is the Markov chain approximation method developed by [Kushner and Dupuis \(2001\)](#), and the specific algorithm given here builds on the publicly available code and algorithm provided by [Alvarez, Argente, and Lippi \(2020\)](#). An online repository with the code, files, and results for our numerical analysis can be found at <https://github.com/mbarnet0/Covid>.

We start by solving the planner's problem for the post-resolution period. Once the vaccine and cure are developed, all concerns about the pandemic drop out, labor is supplied inelastically, and so the planner's problem is present-discounted value of lifetime expected utility given by

$$\hat{V}(N_T, z_T) = E_0 \left[\int_T^\infty e^{-\kappa(t-T)} \kappa \log(C_t) dt \middle| \theta \right]$$

From this, the HJB equation and its solution are trivially derived as

$$\hat{V}(N_T, z_T) = \log(\bar{A}) + \log N_T + \frac{\kappa}{\kappa + 1} z_T$$

With this solution, we can construct the pre-resolution period planner's problem and corresponding HJB equation. The planner's problem is given by

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

Incorporating the solution for $\hat{V}(N_T, z_T)$, the full HJB equation is given by

$$\begin{aligned} (\kappa + \lambda)V = & \max_{q_t} \min_{\tilde{\pi}(\theta)} \kappa \log(\bar{A}N_t) + \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) + \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \\ & + \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \{ (V_i - V_s) \beta s_t i_t (1 - \zeta q_t)^2 + V_s s_t i_t \delta_t - V_i i_t [\gamma - \delta_t i_t] - V_N \delta_t i_t N_t + V_z (-\alpha q_t - z_t) \} \\ & + \frac{1}{2} i_t^2 N_t^2 \sigma_d^2 V_{NN} + \frac{1}{2} \sigma_z^2 V_{zz} + \frac{1}{2} [V_{ss} (\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2 + V_{ii} (\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)] \\ & - V_{si} [\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t (1 - i_t) i_t^2] - V_{sN} s_t i_t^2 \sigma_d^2 - i_t^2 (i_t - 1) \sigma_d^2 V_{iN} + \lambda \log(\bar{A}N_t) + \lambda \frac{\kappa}{\kappa + 1} z_t \end{aligned}$$

Applying the analytical simplification for the value function described in the text, given by $V(s_t, i_t, N_t, z_t; \theta) = \log(\bar{A}) + \log N_t + \frac{\kappa}{\kappa+1} z_t + v(s_t, i_t; \theta)$, allows us to simplify the HJB

equation for the planner's problem to

$$\begin{aligned}
(\kappa + \lambda)v(s_t, i_t) = & \max_{q_t} \min_{\tilde{\pi}(\theta)} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) + \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \\
& + \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \left\{ -\frac{\alpha\kappa}{\kappa + 1} q_t - \delta_t i_t + v_i \beta s_t i_t (1 - \zeta q_t)^2 - v_s \beta s_t i_t (1 - \zeta q_t)^2 + v_s s_t i_t \delta_t - v_i i_t [\gamma - \delta_t i_t] \right\} \\
& + \frac{1}{2} [v_{ss}(\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2 + v_{ii}(\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)] - v_{si} [\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t (1 - i_t) i_t^2] - \frac{1}{2} i_t^2 \sigma_d^2
\end{aligned}$$

To solve the model computationally, we first solve the minimization component of the model. Taking the FOC with respect to $\tilde{\pi}(\theta)$, we find that

$$\tilde{\pi}(s_t, i_t; \theta) = \frac{\pi(\theta) \exp(-\frac{1}{\xi_a} \{ -\frac{\alpha(\theta)\kappa}{\kappa+1} q_t + \delta_t(\theta)(-i_t + v_s s_t i_t + v_i i_t^2) + \beta(\theta) s_t i_t (1 - \zeta(\theta) q_t)^2 (v_i - v_s) \})}{\sum \pi(\theta) \exp(-\frac{1}{\xi_a} \{ -\frac{\alpha(\theta)\kappa}{\kappa+1} q_t + \delta_t(\theta)(-i_t + v_s s_t i_t + v_i i_t^2) + \beta(\theta) s_t i_t (1 - \zeta(\theta) q_t)^2 (v_i - v_s) \})}$$

Note that $\tilde{\pi}(s_t, i_t; \theta)$ is a function of the optimal choice of q_t , which we solve for from the maximization component of the model. Conversely, the FOC for q_t will result in an optimal value that is a function of the optimal $\tilde{\pi}(s_t, i_t; \theta)$. However, rather than solve this nonlinear system jointly with the PDE for the value function, we use a conditional linearity approach as follows. First, we solve for $\tilde{\pi}(s_t, i_t; \theta)$ as a function of q_t , where we use a pre-determined value for q_t based on an initial guess or the previous value determined in our iterative algorithm. Second, we solve for q_t based on the $\tilde{\pi}(s_t, i_t; \theta)$ we derived as a function of a previous q_t value. We then iterate by solving for $\tilde{\pi}(s_t, i_t; \theta)$ as a function of this updated q_t . The iterative procedure continues until the solutions for $\tilde{\pi}(s_t, i_t; \theta)$ and q_t are consistent.

To derive the optimal quarantine choice q_t from the maximization component, we define Q as the terms involving quarantine in the HJB equation, which will depend on $\tilde{\pi}(s_t, i_t; \theta)$. The equation $Q(q_t; s_t, i_t)$ is given by

$$Q(q_t; s_t, i_t) = \kappa \log(1 - q_t) + (v_i - v_s) s_t i_t (\tilde{\beta} - 2\tilde{\beta}\tilde{\zeta} q_t + \tilde{\beta}\tilde{\zeta}^2 q_t^2) - \frac{\tilde{\alpha}\kappa}{\kappa + 1} q_t$$

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

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

Taking first and second derivatives of Q with respect to q_t gives

$$\begin{aligned} Q' &= -\kappa(1 - q_t)^{-1} + (v_i - v_s)s_t i_t (-2\widetilde{\beta\zeta} + 2\widetilde{\beta\zeta^2}q_t) - \frac{\tilde{\alpha}\kappa}{\kappa + 1} \\ Q'' &= -\kappa(1 - q_t)^{-2} + (v_i - v_s)s_t i_t 2\widetilde{\beta\zeta^2} \end{aligned}$$

Setting $Q' = 0$ we can derive the interior solution for the optimal choice of q_t as

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

Note that the complementary slackness condition for the first-order condition implies that $q_{max} \geq q_t \geq 0$. When $Q'' < 0$, then the second-order condition is concave and guarantees the optimal solution to q_t is an interior solution satisfying the equation given above. When $Q'' \geq 0$ the second-order condition is convex and so the optimal solution to q_t is a corner solution. Because of the Inada conditions associated with the use of log utility, when $Q'' \geq 0$ occurs the corner solution must be $q_t = 0$. Furthermore, for numerical stability we impose an upper bound constraint $q_t \leq 0.99$ in the algorithm, though this constraint never binds in our converged solution. As Q'' depends on q_t , and we are simultaneously solving for the value function V and the optimal quarantine policy q_t , for each iteration of the algorithm that we outline below, we use the previous guess for q_t to determine whether $Q'' \geq 0$ or $Q'' < 0$, then solve for the new optimal q_t based on convexity or concavity, update our solution for V , and then use that new value of q_t to check Q'' in the following iteration of the algorithm.

To solve the HJB equation, we separate the state space into the relevant cases over which we need to impose distinct and necessary conditions for the Markov chain approximation method to solve the HJB equation. We only solve the model for combinations of the state variables such that $s_t + i_t \leq 1$, thereby ensuring the adding up constraint $s_t + i_t + r_t = 1$ can be satisfied. We begin by first discretizing the values of the continuous-valued state variables in the model as follows:

$$\begin{aligned} s_s &\in \{0, \Delta_s, \dots, 1 - \Delta_s, 1\} \\ i_i &\in \{0, \Delta_i, \dots, 1 - \Delta_i, 1\} \end{aligned}$$

There is an analytical solution for the boundary of the state space where $i_t = 0$. Note

that because $i_t = 0$ there is no possibility of susceptible becoming infected, therefore the optimal choice of $q_t = 0$, and the analytical solution is given by

$$v(s_t, i_t | i_t = 0) = 0$$

The remaining state-space regions for which we need to solve for a numerical solution to the PDE are given as follows. We first loop over values of s_s , and check whether $s_s = 0$ or $0 < s_s < 1$. For each of these cases, we check whether $0 < i_i < i_{max}$ or $i_i = i_{max}$ where i_{max} is the largest value of i_i in our discretized state space such that $s_s + i_i \leq 1$. Depending on which case we are in, we use a different set of Markov chain Approximation conditions. A visual construction of these conditions in an If-Then code structure is given as follows:

Table A.1: If-Then State-Variable Conditions

s state-variable condition	i state-variable condition
if $s_s = 0$	if $0 < i_i < i_{max}$ else if $i_i = i_{max}$
else if $0 < s_s < 1$	if $0 < i_i < i_{max}$ else if $i_i = i_{max}$

The Markov chain approximation is based on an upwind finite difference scheme to determine transition probabilities for each state at a given point in the state space. Starting from the general HJB equation and replacing value function derivatives with the relevant finite differences, subtracting κv from each side, multiplying by dt , and adding V to each side, we are able to rearrange to get common terms so that our expression to use in our numerical algorithm is given by

$$\begin{aligned}
(\kappa + \lambda)v(s_t, i_t) &= \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - \frac{1}{2}i_t^2\sigma_d^2 - v_{si}[\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t(1 - i_t)i_t^2] \\
&+ \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \\
&+ \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \left\{ \frac{v(s, i + \Delta_i) - v(s, i)}{\Delta_i} \beta s_t i_t (1 - \zeta q_t)^2 - \frac{v(s, i) - v(s - \Delta_s, i)}{\Delta_s} \beta s_t i_t (1 - \zeta q_t)^2 \right. \\
&+ \frac{v(s + \Delta_s, i) - v(s, i)}{\Delta_s} s_t i_t \delta_t - \frac{v(s, i) - v(s, i - \Delta_i)}{\Delta_i} i_t [\gamma - \delta_t i_t] - \frac{\alpha \kappa}{\kappa + 1} q_t - \delta_t i_t \Big\} \\
&+ \frac{1}{2} \frac{v(s + \Delta_s, i) + v(s - \Delta_s, i) - 2v(s, i)}{\Delta_s^2} (\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2 \\
&+ \frac{1}{2} \frac{v(s, i + \Delta_i) + v(s, i - \Delta_i) - 2v(s, i)}{\Delta_i^2} (\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)
\end{aligned}$$

Note that we “hold constant” the cross-partial term $-v_{si}[\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t(1 - i_t)i_t^2]$ by including it in the flow utility at any given iteration in the algorithm and update each of these values with each iteration to find a consistent solution. In addition, we use the conditional linearity approach discussed earlier to solve for $\tilde{\pi}(\theta; s_t, i_t)$ and q_t .

According to the method of [Kushner and Dupuis \(2001\)](#), the condition we need to satisfy for convergence of our algorithm is that

$$dt < (\kappa + \lambda + \frac{\beta s_t i_t (1 - \zeta q_t)^2 + s_t i_t \delta_t}{\Delta s} + \frac{\beta s_t i_t (1 - \zeta q_t)^2 + (\gamma - \delta_t i_t) i_t}{\Delta i} + \frac{(\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2}{\Delta s^2} + \frac{(\sigma_i^2 s_t^2 i_t^2 + (1 - i_t^2) i_t^2 \sigma_d^2)}{\Delta i^2})^{-1}$$

Therefore, we impose

$$dt = 0.95 \times \min_{s \in S, i \in I} (\kappa + \lambda + \frac{\beta s_t i_t + s_t i_t \delta_t}{\Delta s} + \frac{\beta s_t i_t + (\gamma - \delta_t i_t) i_t}{\Delta i} + \frac{(\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2}{\Delta s^2} + \frac{(\sigma_i^2 s_t^2 i_t^2 + (1 - i_t^2) i_t^2 \sigma_d^2)}{\Delta i^2})^{-1}$$

where the min operator is applied over the entire two-dimensional state space for our model. As the above approximation only holds for certain parts of the state-space because of the triangularity imposed by the adding up condition $s_t + i_t \leq 1$, we consider case-by-case the alternative finite-difference approximations imposed in each relevant scenario to adapt this general approximation expression to work in each region of the state-space.

When $s_s = 0, 0 < i_i < i_{max}$:

$$\begin{aligned} v(s_t, i_t) = & \kappa \log(1 - (1 - \phi) i_t) - \tilde{\delta}_t i_t - \frac{1}{2} i_t^2 \sigma_d^2 + \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \\ & + (1 - (\kappa + \lambda) dt) [v(s, i - \Delta_i) - v(s, i)] \left\{ \frac{i_t [\gamma - \tilde{\delta}_t i_t]}{(1 - (\kappa + \lambda) dt) \Delta_i} \frac{dt}{\Delta_i} \right\} \\ & + (1 - (\kappa + \lambda) dt) [v(s, i + \Delta_i) - 2v(s, i) + v(s, i - \Delta_i)] \left\{ \frac{1}{2} \frac{(1 - i_t)^2 i_t^2 \sigma_d^2}{(1 - (\kappa + \lambda) dt) \Delta_i^2} \frac{dt}{\Delta_i^2} \right\} \\ & + (1 - (\kappa + \lambda) dt) v(s, i) \end{aligned}$$

When $s = 0, i = i_{max}$:

$$\begin{aligned}
v(s_t, i_t) &= \kappa \log(1 - (1 - \phi)) - \tilde{\delta}_t - \frac{1}{2} \sigma_d^2 + \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \\
&+ (1 - (\kappa + \nu)dt)[v(s, i - \Delta_i) - v(s, i)] \left\{ \frac{\gamma - \tilde{\delta}_t}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_i} \right\} \\
&+ (1 - (\kappa + \lambda)dt)v(s, i)
\end{aligned}$$

When $0 < s_s < s_{max}, 0 < i_i < i_{max}$:

$$\begin{aligned}
v(s_t, i_t) &= \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - \frac{\tilde{\alpha}\kappa}{\kappa + 1} q_t - \tilde{\delta}_t i_t \\
&- \frac{1}{2} i_t^2 \sigma_d^2 - v_{si} [\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t (1 - i_t) i_t^2] + \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \\
&+ (1 - (\kappa + \lambda)dt)[v(s - \Delta_s, i) - v(s, i)] \left\{ \frac{\beta(1 - \zeta q_t)^2 s_t i_t}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_s} \right\} \\
&+ (1 - (\kappa + \lambda)dt)[v(s + \Delta_s, \max\{i - k \times \Delta_i, 1\}) - v(s, i)] \left\{ \frac{s_t i_t \tilde{\delta}_t}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_s} \right\} \\
&+ (1 - (\kappa + \lambda)dt)[v(s, i - \Delta_i) - v(s, i)] \left\{ \frac{i_t [\gamma - \tilde{\delta}_t i_t]}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_i} \right\} \\
&+ (1 - (\kappa + \lambda)dt)[v(s, i + \Delta_i) - v(s, i)] \left\{ \frac{\beta(1 - \zeta q_t)^2 s_t i_t}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_i} \right\} \\
&+ (1 - (\kappa + \lambda)dt)[v(s + \Delta_s, \max\{i - k \times \Delta_i, 1\}) - 2v(s, i) + v(s - \Delta_s, i)] \left\{ \frac{1}{2} \frac{(\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_s^2} \right\} \\
&+ (1 - (\kappa + \lambda)dt)[v(s, i + \Delta_i) - 2v(s, i) + v(s, i - \Delta_i)] \left\{ \frac{1}{2} \frac{(\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_i^2} \right\} \\
&+ (1 - (\kappa + \lambda)dt)v(s, i)
\end{aligned}$$

In addition we need approximations for $v_i - v_s, v_{si}$. We use

$$\begin{aligned}
v_i - v_s &\approx \frac{v(s, i + \Delta_i) - v(s, i)}{\Delta_i} - \frac{v(s, i) - v(s - \Delta_s, i)}{\Delta_s} \\
v_{si} &\approx \begin{cases} \frac{v(s + \Delta_s, i + \Delta_i) - v(s + \Delta_s, i - \Delta_i) - v(s - \Delta_s, i + \Delta_i) + v(s - \Delta_s, i - \Delta_i)}{4\Delta_s \Delta_i} & \text{when } s_s + \Delta_s + i_i + \Delta_i \leq 1 \\ \frac{v(s, i + \Delta_i) - v(s, i - \Delta_i) - v(s - \Delta_s, i + \Delta_i) + v(s - \Delta_s, i - \Delta_i)}{2\Delta_s \Delta_i} & \text{otherwise} \end{cases}
\end{aligned}$$

When $0 < s_s < s_{max}, i = i_{max}$:

$$\begin{aligned}
v(s_t, i_t) = & \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - \frac{\tilde{\alpha}\kappa}{\kappa + 1} q_t - \tilde{\delta}_t i_t \\
& - \frac{1}{2} i_t^2 \sigma_d^2 - v_{si} [\sigma_i^2 s_t^2 i_t^2 + \sigma_d^2 s_t (1 - i_t) i_t^2] + \xi_a \sum_{\theta \in \Theta} \tilde{\pi}(\theta) \log \frac{\tilde{\pi}(\theta)}{\pi(\theta)} \\
& + (1 - (\kappa + \lambda)dt) [v(s - \Delta_s, i + k \times \Delta_i) - v(s, i)] \left\{ \frac{\beta(1 - \zeta q_t)^2 s_t i_t}{(1 - (\kappa + \lambda)dt) \Delta_s} \frac{dt}{\Delta_s} \right\} \\
& + (1 - (\kappa + \lambda)dt) [v(s + \Delta_s, i - k \times \Delta_i) - v(s, i - k \times \Delta_i)] \left\{ \frac{s_t i_t \tilde{\delta}_t}{(1 - (\kappa + \lambda)dt) \Delta_s} \frac{dt}{\Delta_s} \right\} \\
& + (1 - (\kappa + \lambda)dt) [v(s, i - \Delta_i) - v(s, i)] \left\{ \frac{i_t [\gamma - \tilde{\delta}_t i_t]}{(1 - (\kappa + \lambda)dt) \Delta_i} \frac{dt}{\Delta_i} \right\} \\
& + (1 - (\kappa + \lambda)dt) [v(s + \Delta_s, i - k \times \Delta_i) - 2v(s, i - k \times \Delta_i) + v(s - \Delta_s, i - k \times \Delta_i)] \\
& \times \left\{ \frac{1}{2} \frac{(\sigma_d^2 + \sigma_i^2) s_t^2 i_t^2}{(1 - (\kappa + \lambda)dt) \Delta_s^2} \frac{dt}{\Delta_s^2} \right\} \\
& + (1 - (\kappa + \lambda)dt) [v(s - \Delta_s, i + k \times \Delta_i) - 2v(s - \Delta_s, i) + v(s - \Delta_s, i - k \times \Delta_i)] \\
& \times \left\{ \frac{1}{2} \frac{(\sigma_i^2 s_t^2 i_t^2 + (1 - i_t)^2 i_t^2 \sigma_d^2)}{(1 - (\kappa + \lambda)dt) \Delta_i^2} \frac{dt}{\Delta_i^2} \right\} \\
& + (1 - (\kappa + \lambda)dt) v(s, i)
\end{aligned}$$

In addition we need approximations for $v_i - v_s, v_{si}$. We use

$$\begin{aligned}
v_i - v_s & \approx \frac{v(s - \Delta_s, i + \Delta_i) - v(s, i)}{\Delta_s} \\
v_{si} & \approx \frac{v(s, i) - v(s, i - \Delta_i) - v(s - \Delta_s, i) + v(s - \Delta_s, i - \Delta_i)}{\Delta_s \Delta_i}
\end{aligned}$$

With the set-up for the Markov chain approximation set, we use the following hyper-parameters and initial guesses in order to compute the numerical solutions used in our analysis. First, we use 650 points along the s_t dimension, $N_s = 650$, and five times as many points along the i_t dimensions, $N_i = N_s \times 5 = 3250$. This leads to a discretization step size along the s_t dimension that is $\Delta_s = \frac{1}{N_s - 1} = 0.0015$ and along the i_t dimension that is $\Delta_i = \frac{1}{N_i - 1} = 0.00031$. The resulting state space includes over 2 million unique (s_t, i_t) points, though this is significantly reduced by the fact that the adding up constraint of $1 = s_t + i_t + r_t$ makes certain points infeasible and therefore unused in our computations. We use two different initial guesses for the value function based on specific cases. For all but one of the individual models used in the “outside the model” uncertainty comparison,

as well as for the uncertainty neutral cases where $\xi_a = \infty$, we use the same initial guess of

$$V_0 = \frac{\kappa \log(1 - (1 - \phi)i_t)}{\kappa + \lambda} - \frac{\kappa \log(1 - s_t i_t)}{\kappa + \lambda}$$

Because the case with $(\mathcal{R}_0, CFR, \zeta, \alpha) = (1.5, 0.02, 0.65, 0.0)$ produces particularly poor derivative approximations near the boundary, we use the solution to the case with $(\mathcal{R}_0, CFR, \zeta, \alpha) = (1.5, 0.02, 0.35, 0.0)$ as the initial guess for more stable results. For the uncertainty averse cases, we use the uncertainty neutral solution with the same prior weighting as the initial value function guess V_0 . This ensures stability in convergence with the additional curvature and decreases the number of iterations required for convergence. Finally, convergence is achieved when the error terms for the policy function and the value function are below given tolerance levels. For the value function error, call it *ValueFunctionError*, we set the tolerance to be $\epsilon_1 = 1\text{e-}12$ and use the same criteria as in [Alvarez et al. \(2020\)](#):

$$ValueFunctionError = \kappa \frac{\|V^n(s, i) - V^{n-1}(s, i)\|_1}{n}$$

where $\|\cdot\|_1$ represents the 1-norm and n denotes the iteration number in the Markov chain approximation algorithm. We add an additional error term based on the policy function, call it *PolicyFunctionError*, in order to ensure convergence. We set this tolerance to be $\epsilon_2 = 5\text{e-}8$ and use the criteria

$$PolicyFunctionError = \frac{\|q_t^n - q_t^{n-1}\|_2}{\|q_t^{n-1}\|_2}$$

where $\|\cdot\|_2$ represents the 2-norm and again n denotes the iteration number in the Markov chain approximation algorithm. The solutions are qualitatively similar for various different hyper-parameters and initial guesses for which convergence is achieved, and any quantitative differences that arise in a small number of cases are minor. For the most extreme “outside the model” cases with the highest mitigation policy outcomes, i.e., the cases with high infection and fatality rates and low quarantine effectiveness, occasionally the iterative scheme will generate poor derivative approximations near the boundary which produces unreliable results for some choices of the hyper-parameters. We tested numerous hyper-parameter choices to verify that results are consistent across cases where this does not occur.

Once we have model solutions using our algorithm, results are generated using simulations based on the model solutions where model priors and underlying processes are chosen to fit

the scenarios described in the main text. The simulations set shocks to zero, while using solutions assuming there are shocks in the model. Initial values for the simulations are chosen by assuming there is a small initial population of infected individuals, as well as a small initial value of immune/recovered individuals. The first assumption is required to provide a small but rapidly growing level of infection in the model as was observed in the US, and the second is in line with [Alvarez et al. \(2020\)](#) and consistent with some recent research that some people may have at least partial immunity¹⁸, and ensures stability of the simulations that avoid numerical inaccuracies near boundary values as well. The initial values we use are given by $s_0 = 1 - i_0 - r_0$, $i_0 = 0.02$, $r_0 = 0.03$, and $d_0 = 0$.

Appendix E Alternative Prior Weight Scenarios

We provide results for three additional sets of prior beliefs scenarios over the models which we label as follows: (i) “under-estimated,” (ii) “over-estimated”, and (iii) “split-estimated.” Each case highlights important scenarios the planner could confront, and how uncertainty influences policy choices and model outcomes in those different scenarios. For each scenario we show the following results from the uncertainty neutral and the uncertainty averse solutions: Figures [A.1](#), [A.4](#), and [A.7](#) show the planner’s optimal quarantine and resulting infections and deaths; Figures [A.2](#), [A.5](#), and [A.8](#) show how the planner’s beliefs over \mathcal{R}_0 , the CFR, the mitigation policy losses, and the mitigation policy effectiveness evolve correspondingly; Figures [A.3](#), [A.6](#), and [A.9](#) show the uncertainty averse planner’s distorted probabilities over each model and the detection error probability values for the uncertainty-adjusted solutions.

We first consider the “under-estimated” scenario where the planner assigns a weight of 50% to the model with $\mathcal{R}_0 = 1.5$, initial CFR= 0.005, $\zeta = 0.65$, and $\alpha = 0.0$, with the remaining weight distributed equally to the remaining models. In this case, Figure [A.1](#) shows that the uncertainty neutral optimal policy response is quite modest, with quarantine measures starting at zero, peaking around 45%, and finishing after about 12 weeks. This response results in a relatively high number of deaths (almost 2%). The uncertainty averse response significantly increases the initial quarantine level (near 20%), the peak quarantine level (over 60%) and the duration of the quarantine measures (about 20 weeks), drastically reducing the number of deaths. Figures [A.2](#) and [A.3](#) show that this policymaker substantially decreases weight on the best-case model, transferring much of that weight to the two most-severe models as infections and deaths rise. This leads to implied \mathcal{R}_0 , CFR , and α that are much higher than originally assumed, and an implied ζ that is much lower than assumed

¹⁸See articles in [Science](#), [Nature](#), and [medRxiv](#) for examples of recent research discussing this.

by the prior.¹⁹ The policymaker reassigns weight to the worst-case outcomes because as the pandemic initially spreads quickly through the population the large number of infected means that underestimating the pandemic becomes even more costly and the planner therefore takes extreme measures to contain the pandemic going forward, approaching policy close to the equal-weighted case in the main text. Thus, in a case where the planner underestimates the severity and economic costs of the pandemic, incorporating uncertainty substantially increases the quarantine policy choice, increasing as the pandemic runs its course, and only shifting back to the prior after a sustained period of low infections.

Next we consider the “over-estimated” scenario where the planner assigns a weight of 50% to the model with $\mathcal{R}_0 = 4.5$, initial CFR= 0.02, $\zeta = 0.35$, and $\alpha = 0.8$, with the remaining weight distributed equally to the remaining models. In this case, Figure A.4 shows a higher quarantine policy for the uncertainty neutral solution, with quarantine measures reaching nearly 50% at around 8-10 weeks. The deaths resulting from this response are modest compared to the other scenarios. However, the uncertainty averse response has only a modest impact on the quarantine response, peaking between 55% and 60%, and returning to zero just before 20 weeks. Figures A.5 and A.6 show that the policymaker maintains high weight on the most-severe model, and has a somewhat higher implied \mathcal{R}_0 , CFR , and α , and a lower implied ζ . However, given the already high initial weight on the worst-case model, and the fact that the realized outcomes are in fact less severe than what the initial prior weighting would have anticipated, the potential magnitude of probability weights is somewhat limited.

Finally, we consider the “split-estimated” scenario where the planner assigns a weight of 26.67% to the model with $\mathcal{R}_0 = 4.5$, initial CFR= 0.02, $\zeta = 0.65$, and $\alpha = 0.0$, a weight of 26.67% to the model with $\mathcal{R}_0 = 1.5$, initial CFR= 0.005, $\zeta = 0.35$, and $\alpha = 0.8$ with the remaining weight distributed equally to the remaining models. The quarantine, infection, and death outcomes are shown for this case in Figure A.7. Given the identical weighted-average parameter values resulting from this case as in the equal-weighted case, the uncertainty neutral results are the same as in that baseline setting. In terms of the uncertainty averse response, we see again that there is a noticeable impact on quarantine measures. However, the effect is more modest in this setting. Figures A.8 and A.9 show that the policymaker has a fairly high weight on the most severe model, but still maintains a relatively high weight on the least severe model. Thus, the structure of the prior and not just the weighted-value outcomes of the prior, are shown to have some impact. In addition, the implied \mathcal{R}_0 , CFR , and α , and ζ values show movements similar to the equal-weighted

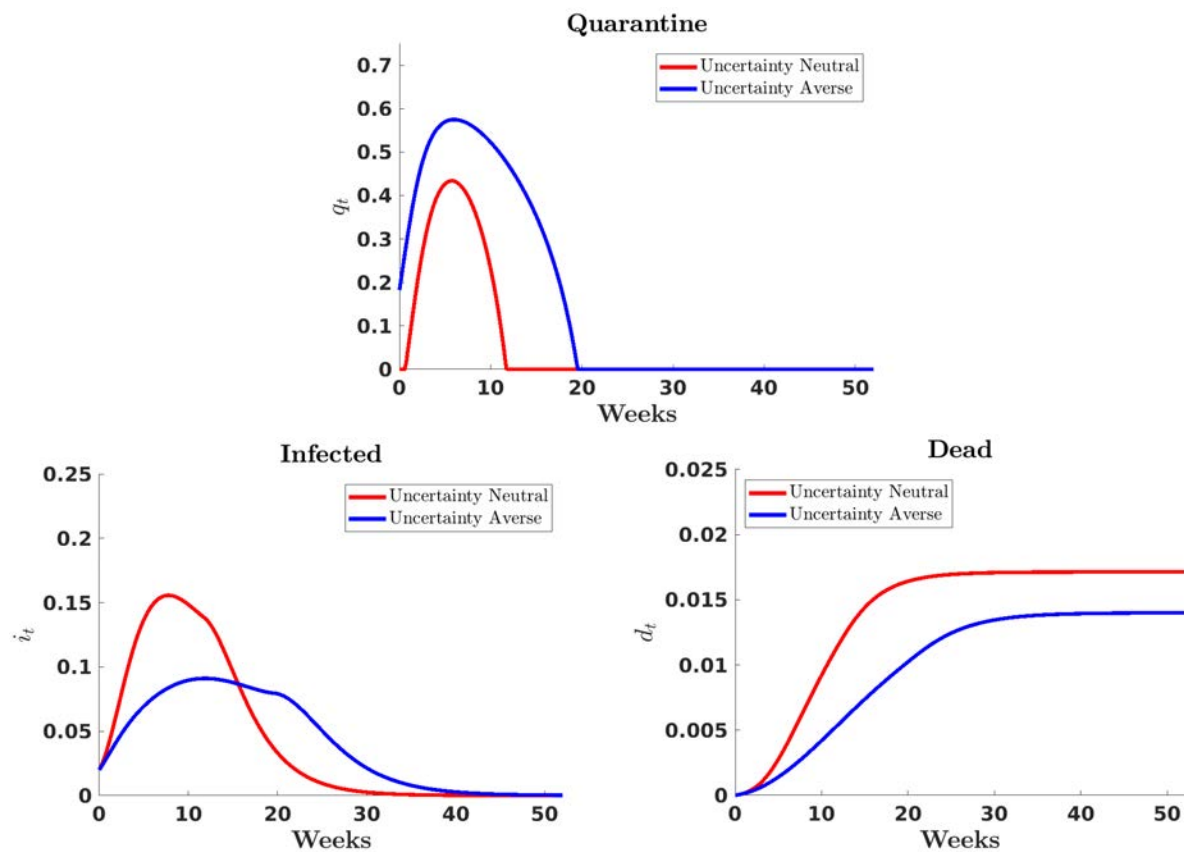
¹⁹In the uncertainty neutral case, the belief over the CFR increases because the death rate is state-dependent and there is a significant increase in the number of infections.

baseline, but are again less significant given this alternative weighting.

The detection error probabilities for each case are shown in Figures A.3, A.6, and A.9. We can see that for each case, we have adjusted the value of ξ_a for each scenario to keep the minimum value roughly consistent across settings. The “under-estimated” case uses an uncertainty parameter value of $\xi_a = 0.0042$, the “over-estimated” cases uses an uncertainty parameter value of $\xi_a = 0.002$, and the “split-estimated” case uses an uncertainty parameter value of $\xi_a = 0.00745$. Each case, apart from the “over-estimated” setting, has a minimum value slightly above 10%. There are a few key differences to note across cases. First, the “under-estimated” and “over-estimated” cases never converge back to the assumed prior weight of 50%, because the model comparisons between uncertainty neutral and uncertainty averse models assume the “under-estimated” and “over-estimated” priors but use the equal-weighted parameter values to generate the model outcomes. As a result, there is always some detection error probability because the assumed model and the data generating model are in fact different. Also, the “split-estimated” converges much more slowly back to the 50% assumed prior weight value because of the more severe divergence in the weighting of the discrete set of model possibilities. Finally, note that the detection error probability does not go below 12%. This is due to the fact that the initial weight on worst-case model is already substantial, limiting potential for probability weight distortions. Combined with the fact that the realized outcomes are in fact less severe than what the initial prior weighting would have anticipated, this limits the magnitude of potential model detection error.

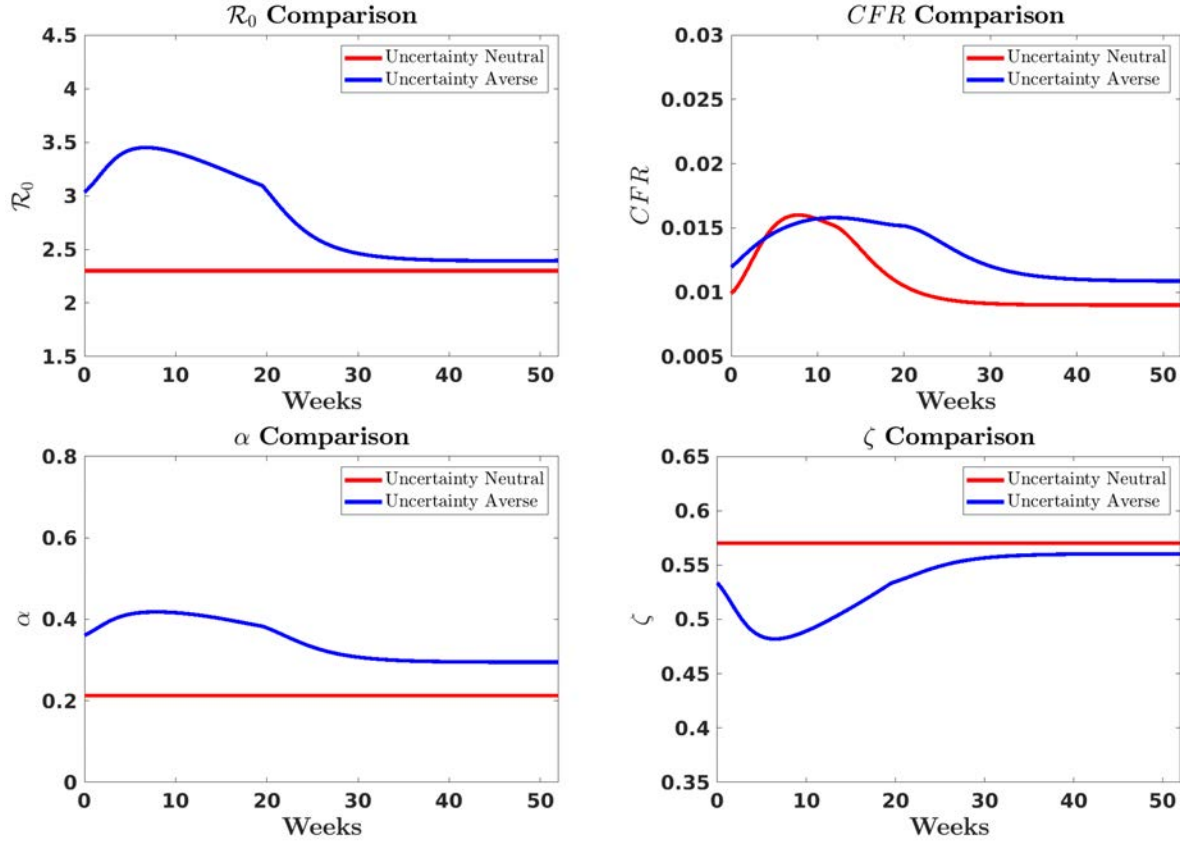
These results demonstrate the asymmetric impact of uncertainty depending on the policymakers initial beliefs about the severity and economic costs of the pandemic. In each case, uncertainty leads to higher and more persistent policy responses than the uncertainty neutral counterpart, reinforcing our baseline model results. However, the results from these cases show that the impact of incorporating uncertainty is greatest when we initially underestimate the severity and economic costs of the pandemic. This is because the pandemic evolution is the most negatively different from the assumed prior expectation, justifying and allowing for the largest distortions to the prior model weights towards the worst-case model. The smallest relative increase comes when we initially overestimate the severity and economic costs of the pandemic because the pandemic evolution is the least negatively different from the assumed prior expectation. Though there is additional weight placed on the worst-case model because of ambiguity aversion when we overestimate the severity and economic costs of the pandemic, this difference is modest because the planner does not entertain the same relative increase in concerns about the worst-case model as when the pandemic severity is underestimated, leading to the smallest distortions in relative terms.

Figure A.1: Outcomes With and Without Uncertainty, “Under-Estimated” Prior



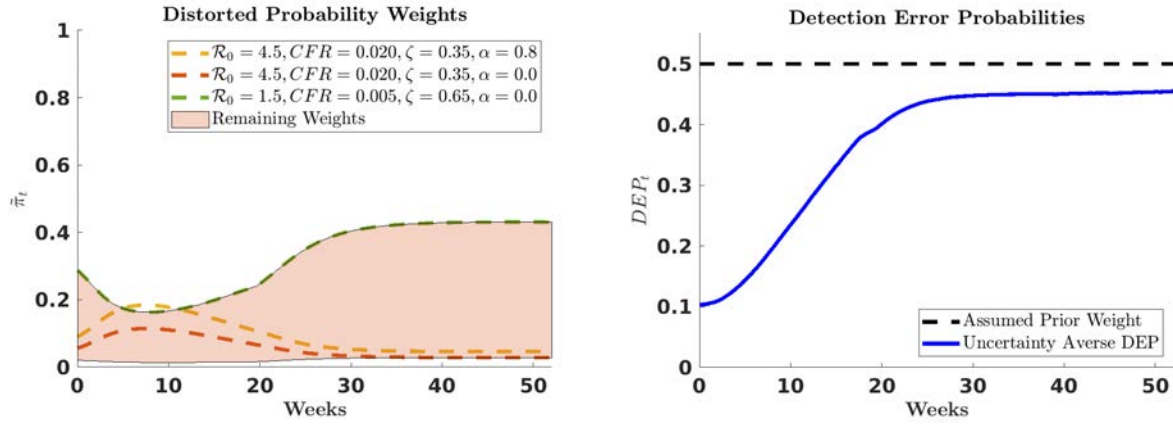
Notes: These figures show the fraction of the population quarantined (top), infected (bottom left), and dead (bottom right) under (i) the uncertainty neutral model response (red) and the uncertainty averse model response (blue).

Figure A.2: Distorted Model Parameters: $\mathcal{R}_0, CFR, \alpha, \zeta$



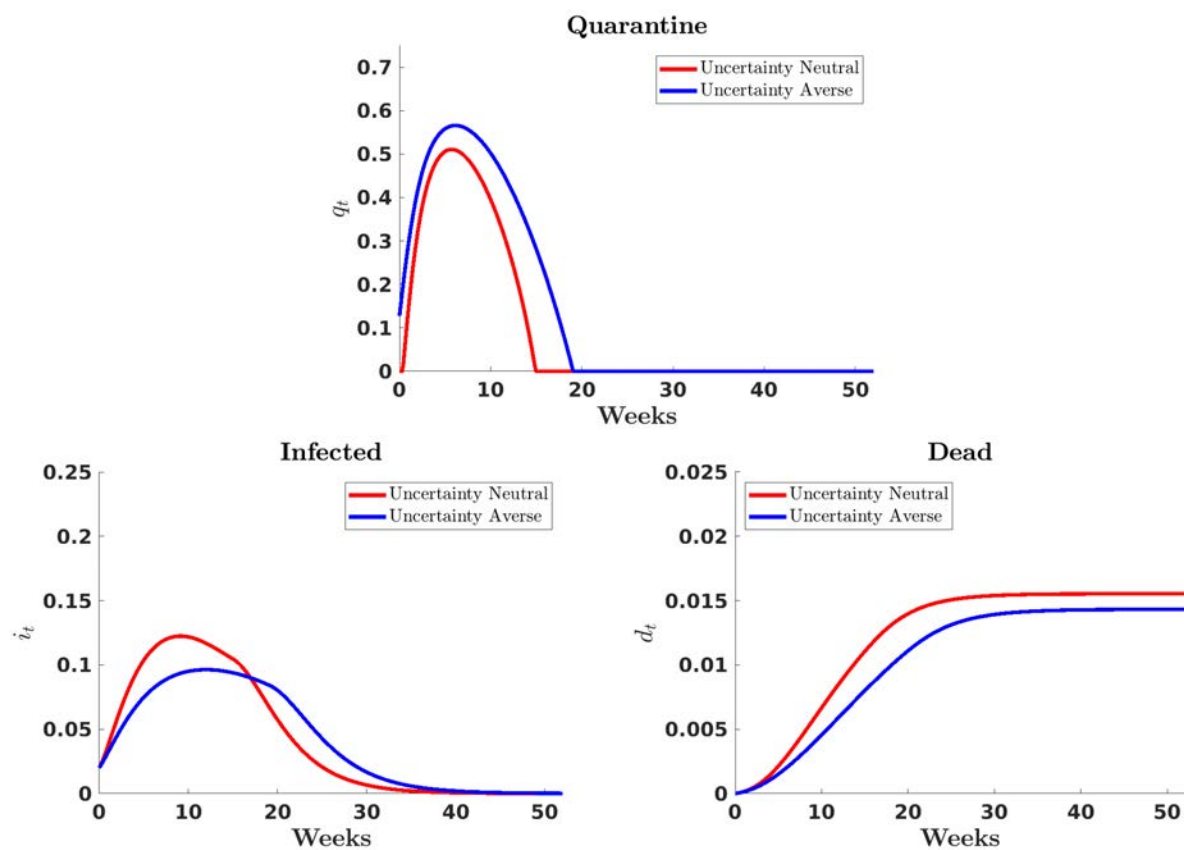
Notes: The figures show the implied \mathcal{R}_0 (top left), CFR (top right), α (bottom left), and ζ (bottom right) for the uncertainty neutral case (red lines) and the uncertainty averse case (blue lines). Each case uses the simulation path resulting from its own optimal policy.

Figure A.3: Uncertainty Aversion Distorted Model Probabilities and Detection Error Probability Comparison



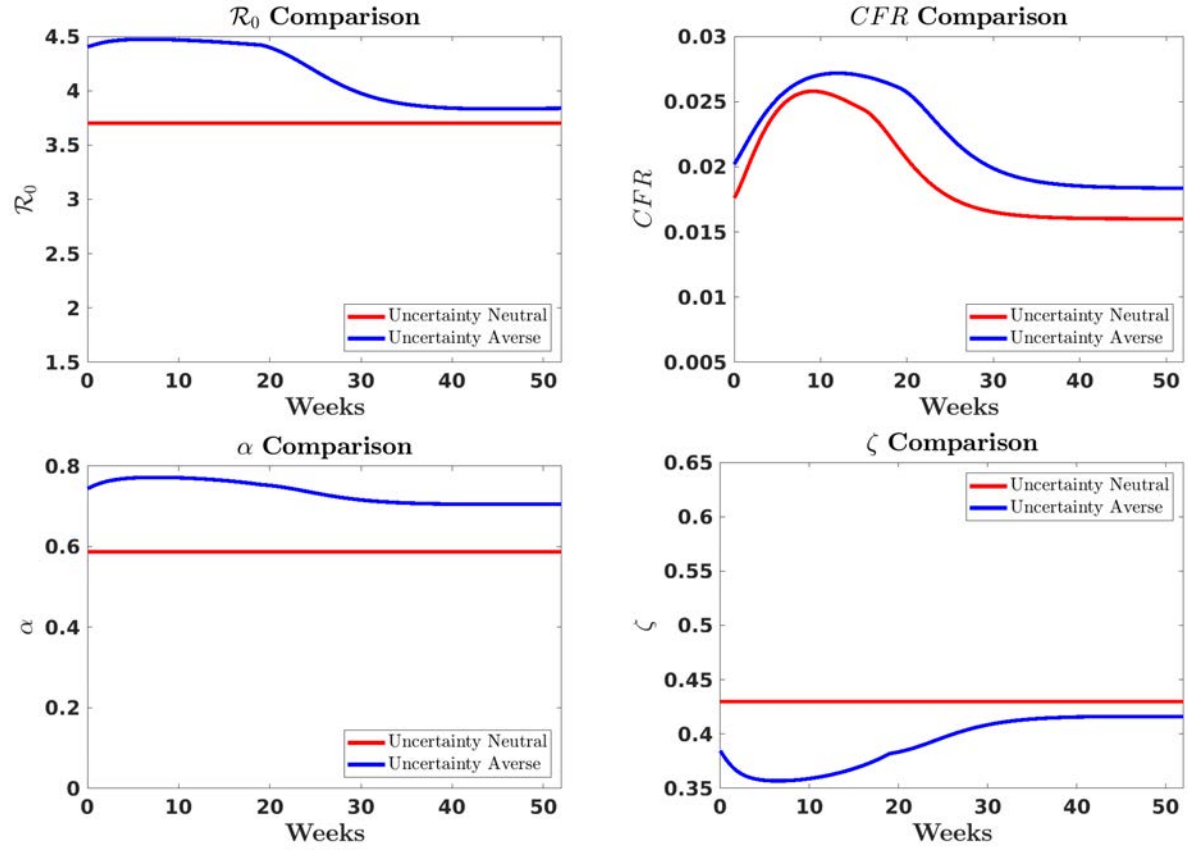
Notes: The left panel shows the distorted probability weights for the uncertainty averse planner. The red dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The yellow dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the highest economic cost of quarantine ($\alpha = 0.8$). The green dashed line is the model with the least severe epidemiological scenario ($CFR = 0.005, \mathcal{R}_0 = 1.5$), the most effective pandemic mitigation scenario ($\zeta = 0.65$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The red shaded area shows the range of distorted probability weights for the remaining models. The right panel 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).

Figure A.4: Outcomes With and Without Uncertainty, “Over-Estimated” 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 the uncertainty averse model response (blue).

Figure A.5: Distorted Model Parameters: $\mathcal{R}_0, CFR, \alpha, \zeta$



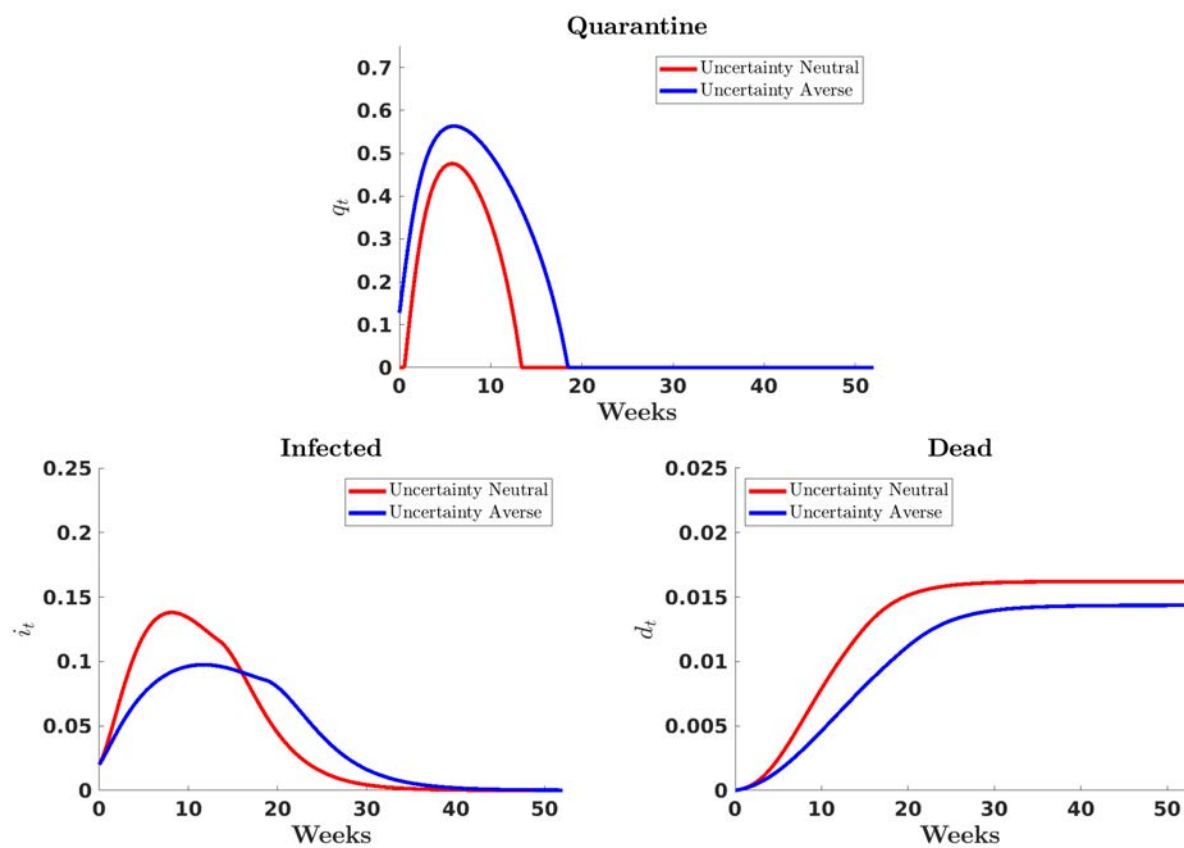
Notes: The figures show the implied \mathcal{R}_0 (top left), CFR (top right), α (bottom left), and ζ (bottom right) for the uncertainty neutral case (red lines) and the uncertainty averse case (blue lines). Each case uses the simulation path resulting from its own optimal policy.

Figure A.6: Uncertainty Aversion Distorted Model Probabilities and Detection Error Probability Comparison



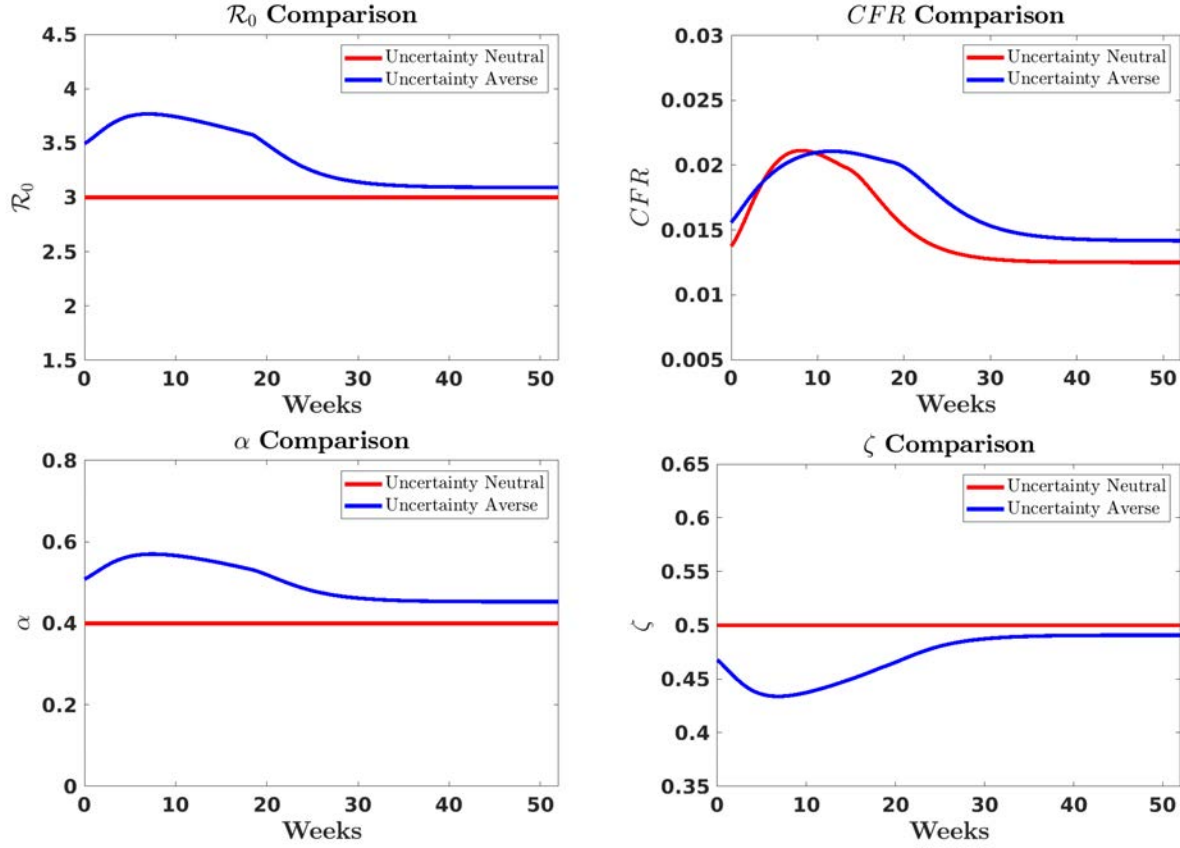
Notes: The left panel shows the distorted probability weights for the uncertainty averse planner. The dashed lines are the probability weights on the most severe epidemiological scenario ($CFR = 0.02$, $\mathcal{R}_0 = 4.5$) with the least effective pandemic mitigation scenario ($\zeta = 0.35$). The 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. The right panel 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).

Figure A.7: Outcomes With and Without Uncertainty, “Split-Estimated” 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 the uncertainty averse model response (blue).

Figure A.8: Distorted Model Parameters: $\mathcal{R}_0, CFR, \alpha, \zeta$



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.

Figure A.9: Uncertainty Aversion Distorted Model Probabilities and Detection Error Probability Comparison



Notes: The left panel shows the distorted probability weights for the uncertainty averse planner. The red dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The yellow dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the highest economic cost of quarantine ($\alpha = 0.8$). The green dashed line is the model with the least severe epidemiological scenario ($CFR = 0.005, \mathcal{R}_0 = 1.5$), the most effective pandemic mitigation scenario ($\zeta = 0.65$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The red shaded area shows the range of distorted probability weights for the remaining models. The right panel 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).

References

- Abel, A. B. and S. Panageas (2020). Optimal management of a pandemic in the short run and the long run. *NBER Working Paper* (w27742).
- Acemoglu, D., V. Chernozhukov, I. Werning, and M. Whinston (2020). A multi-risk sir model with optimally targeted lockdown. Technical report, National Bureau of Economic Research.
- Ai, H., R. Bansal, H. Guo, and A. Yaron (2019). Identifying preference for early resolution from asset prices. *Working Paper*.
- Alvarez, F. E., D. Argente, and F. Lippi (2020). A simple planning problem for covid-19 lockdown. Technical report, National Bureau of Economic Research.
- Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association* 1(1), 68–123.
- Anscombe, F. J. and R. J. Aumann (1963). A definition of subjective probability. *Annals of Mathematical Statistics* 34(1), 199–205.
- Arrow, K. J. (1951). Alternative approaches to the theory of choice in risk-taking situations. *Econometrica: Journal of the Econometric Society*, 404–437.
- Arroyo-Marioli, F., F. Bullano, S. Kucinskas, and C. Rondón-Moreno (2021). Tracking r of covid-19: A new real-time estimation using the kalman filter. *PloS one* 16(1), e0244474.
- Atkeson, A. (2020). How deadly is covid-19? understanding the difficulties with estimation of its fatality rate. Technical report, National Bureau of Economic Research.
- Atkeson, A., K. Kopecky, and T. Zha (2021). Behavior and the transmission of covid-19. *AEA Papers and Proceedings*.
- Avery, C., W. Bossert, A. Clark, G. Ellison, and S. F. Ellison (2020). Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists. Technical report, National Bureau of Economic Research.
- Baek, J., V. F. Farias, A. Georgescu, R. Levi, T. Peng, D. Sinha, J. Wilde, and A. Zheng (2021). The limits to learning a diffusion model. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 130–131.
- Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics* 131(4), 1593–1636.
- Baker, S. R., N. Bloom, S. J. Davis, and S. J. Terry (2020). Covid-induced economic uncertainty. Technical report, National Bureau of Economic Research.
- Barnett, M., W. Brock, and L. P. Hansen (2020). Pricing uncertainty induced by climate change. *The Review of Financial Studies* 33(3), 1024–1066.
- Barnett, M., W. Brock, and L. P. Hansen (2021). Climate Change Uncertainty Spillover in the Macroeconomy. *NBER Macroeconomics Annual* 36(1), 387–395.
- Barrot, J.-N., B. Grassi, and J. Sauvagnat (2020). Sectoral effects of social distancing. *Working Paper*.
- Berger, L., N. Berger, V. Bosetti, I. Gilboa, L. P. Hansen, C. Jarvis, M. Marinacci, and R. Smith (2020). 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* (2020-95).

- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica* 77(3), 623–685.
- Borovička, J. (2020). Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. *Journal of Political Economy* 128(1), 206–251.
- Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association* 71(356), 791–799.
- Cagetti, M., L. P. Hansen, T. Sargent, and N. Williams (2002). Robustness and pricing with uncertain growth. *The Review of Financial Studies* 15(2), 363–404.
- Chater, N. (2020). Facing up to the uncertainties of covid-19. *Nature Human Behaviour*, 1–1.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics* 23, 493–507.
- Cogley, T., R. Colacito, L. P. Hansen, and T. J. Sargent (2008). Robustness and us monetary policy experimentation. *Journal of Money, Credit and Banking* 40(8), 1599–1623.
- Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020). The macroeconomics of epidemics. Technical report, National Bureau of Economic Research.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 643–669.
- Garlappi, L., R. Uppal, and T. Wang (2007). Portfolio selection with parameter and model uncertainty: A multi-prior approach. *The Review of Financial Studies* 20(1), 41–81.
- Gilboa, I., S. Minardi, and L. Samuelson (2020). Theories and cases in decisions under uncertainty. *Games and Economic Behavior*.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Hansen, L. and T. J. Sargent (2001). Robust control and model uncertainty. *American Economic Review* 91(2), 60–66.
- Hansen, L. P. and J. Miao (2018). Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proceedings of the National Academy of Sciences* 115(37), 9163–9168.
- Hansen, L. P. and T. J. Sargent (2010). Fragile beliefs and the price of uncertainty. *Quantitative Economics* 1(1), 129–162.
- Hansen, L. P. and T. J. Sargent (2011). Robustness and Ambiguity in Continuous Time. *Journal of Economic Theory* 146(3), 1195–1223.
- Hansen, L. P., T. J. Sargent, G. Turmuhambetova, and N. Williams (2006). Robust Control and Model Misspecification. *Journal of Economic Theory* 128(1), 45–90.
- Hansen, L. P., T. J. Sargent, and N. E. Wang (2002). Robust permanent income and pricing with filtering. *Macroeconomic dynamics* 6(1), 40–84.
- Izhakian, Y. (2020). A theoretical foundation of ambiguity measurement. *Journal of Economic Theory* 187, 105001.
- Izhakian, Y. and D. Yermack (2017). Risk, ambiguity, and the exercise of employee stock options. *Journal of Financial Economics* 124(1), 65–85.
- Jones, C., T. Philippon, and V. Venkateswaran (2021). Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies* 34(11), 5188–5223.

- Kaplan, G., B. Moll, and G. Violante (2020). Pandemics according to hank. Technical report, Working Paper.
- Knight, F. H. (1921). Risk, uncertainty and profit, 1921. *Library of Economics and Liberty*.
- Korolev, I. (2020). Identification and Estimation of the SEIRD Epidemic Model for COVID-19. *Working Paper*.
- Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020). Scarring body and mind: the long-term belief-scarring effects of covid-19. Technical report, National Bureau of Economic Research.
- Kushner, H. J. and P. Dupuis (2001). *Numerical Methods for Stochastic Control Problems in Continuous Time*, Volume 24. Springer Science & Business Media.
- Lemoine, D. M. and C. P. Traeger (2012, July). Tipping Points and Ambiguity in the Economics of Climate Change. Working Paper 18230, National Bureau of Economic Research.
- Li, X., B. Nezami Narajabad, and T. P. Temzelides (2014). Robust dynamic optimal taxation and environmental externalities.
- Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6), 1447–1498.
- Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of Financial Studies* 17(4), 951–983.
- Manski, C. F. and F. Molinari (2020). Estimating the covid-19 infection rate: Anatomy of an inference problem. Technical report, National Bureau of Economic Research.
- Newman, C. M. and B. Stuck (1979). Chernoff bounds for discriminating between two markov processes. *Stochastics: An International Journal of Probability and Stochastic Processes* 2(1-4), 139–153.
- Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.
- Spychalski, P., A. Błażyńska-Spychalska, and J. Kobiela (2020). Estimating case fatality rates of covid-19. *The Lancet Infectious Diseases*.
- Stock, J. H. (2020). Data gaps and the policy response to the novel coronavirus. Technical report, National Bureau of Economic Research.
- Wald, A. (1950). Statistical decision functions.
- 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.
- Yamamoto, N., B. Jiang, and H. Wang (2021). Quantifying compliance with covid-19 mitigation policies in the us: A mathematical modeling study. *Infectious Disease Modelling* 6, 503–513.