EPIDEMIC RESPONSES UNDER UNCERTAINTY

Michael Barnett
Greg Buchak
Constantine Yannelis

Epidemic Responses Under Uncertainty
Michael Barnett, Greg Buchak, and Constantine Yannelis
NBER Working Paper No. 27289
May 2020, Revised April 2021
JEL No. E1,H0,I1

ABSTRACT

We examine how policymakers react to a pandemic with uncertainty regarding key epidemiological parameters by embedding a macroeconomic SIR model in a robust control framework. We find that optimal policy under uncertainty generates optimal mitigation responses that are asymmetric with respect to the initial estimate of the pandemic's severity. When underestimating the severity, the robust control approach leads to a harsher quarantine, closer to the true optimal level, compared to a naive approach. When overestimating, the planner initially implements a policy similar to the true optimal policy but fails to relax it as the pandemic abates.

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

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

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

1   Introduction

The rapid spread of COVID-19 in 2020 was accompanied by a vigorous debate over the costs and benefits of actions taken to mitigate the spread of the pandemic. This debate occurred surrounded by significant uncertainty about key parameters relating to the costs of the new virus, including death rates, infection rates and the economic costs of policies such as shuttering businesses and issuing shelter-in-place orders (Chater, 2020). Many policymakers, academics and commentators in the media suggested that this parameter uncertainty argued for a laxer quarantine and lockdown response.[1] It is not obvious that this is true, and this paper formally explores how optimal pandemic mitigation policies change when faced with significant uncertainty.

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). For $\mathcal{R}_0$, estimates ranged from 1.5 to 12 (Korolev, 2020). Many 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 rage of estimates and highlight how a lack of testing and the presence of many asymptomatic carriers made measurement difficult. To highlight the wide variation in these estimates, Figure 1 shows CFRs and $\mathcal{R}_0$ across many countries and US states.

Knowledge of the CFR and $\mathcal{R}_0$ is critical in informing how to balance the economic costs of mitigation with the public health costs of the disease, and policymakers found themselves confronting critical decisions with incomplete knowledge. Fortunately, economic theory suggests a way forward when confronted with parameter uncertainty. In particular, Wald (1950), Gilboa and Schmeidler (1989), and more recently Hansen and Sargent (2001)

---

[1]For example, New York Mayor De Blasio in a March 9 press conference saying "*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 Ioannides 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.*"

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

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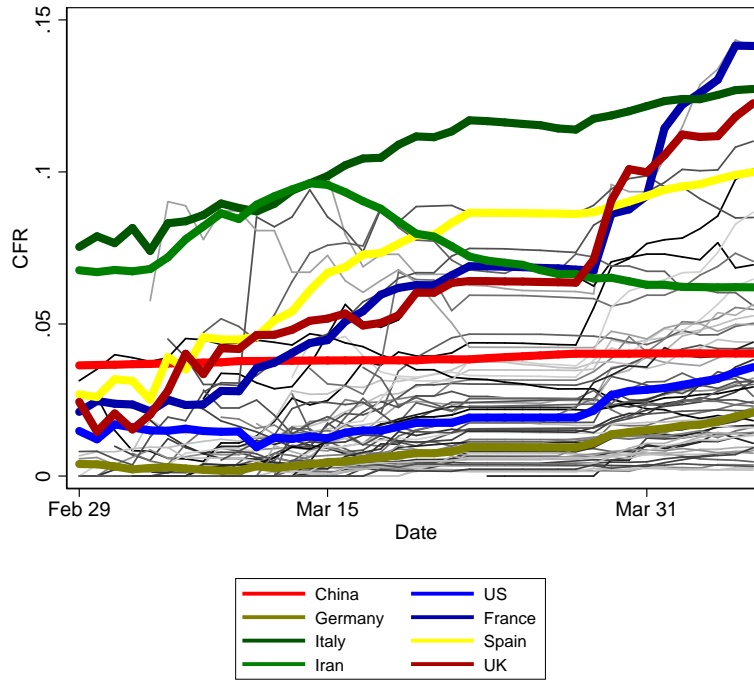
Panel B: Estimated $\mathcal{R}_0$ for Countries and States
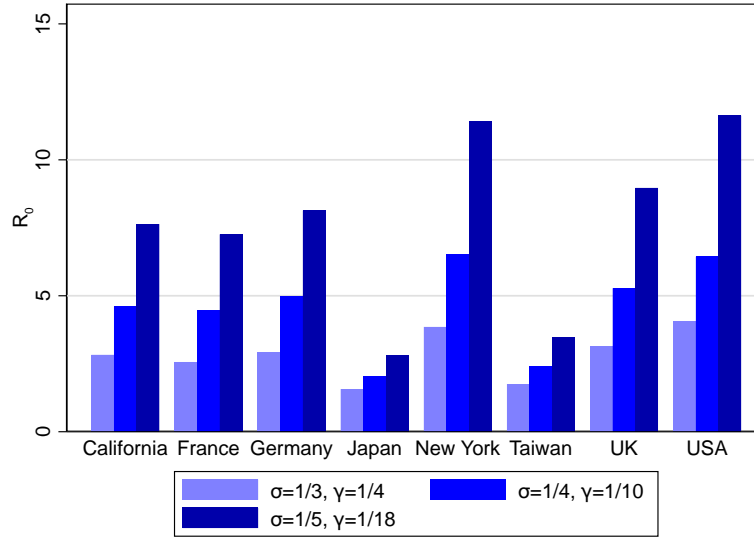
Panel A shows estimated CFRs for all countries with more than 1,000 cases and 100 deaths. Panel B shows estimated $\mathcal{R}_0$ across countries and states given different incubation periods (σ) and durations (γ). Source: European Centre for Disease Prevention and Control and Korolev (2020).

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 10%, however, while towards the extreme 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 to inform optimal policy decisions.

We begin with a simple macroeconomic model with an epidemic. The epidemic spreads according to a standard Susceptible-Infectious-Recovered (SIR) process employed in epidemiology and more recently in macroeconomic models.[2] Infected individuals are less productive, spread the disease, and may die. We allow policymakers to impose a quarantine of varying strictness. The quarantine slows the disease's spread by removing individuals from the working population. The optimal policy balances spread reduction against temporary lost productivity, and depends critically on the CFR and $\mathcal{R}_0$ in the underlying SIR model. We first show the wide range of optimal responses across different underlying parameters. These 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.[3] 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 from the

---

[2]SIR models are standard tools in epidemiology used to model the spread of infectious diseases. The epidemiological SIR model computes the theoretical number of people infected with a contagious disease in a closed population over time. The models have three key elements: S is the number of susceptible, I is the number of infectious, and R is the number of recovered, deceased, or immune individuals. 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.

[3]A large 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).

policymaker. Facing this ambiguity, the policymaker then makes decisions considering worst-case outcomes across alternate parameter settings, where the cost of considering any particular set of parameters depends on how far the parameters lie from prior beliefs. We calibrate the model to match the US economy, and explore how uncertainty influences optimal quarantine policy.

We find that greater parameter uncertainty pushes the optimizing policymaker towards stronger and more persistent quarantine. When the policy maker's initial prior underestimates the severity of the disease, without uncertainty aversion the quarantine policy is quite low and ends relatively quickly, even though infections and deaths reach very high levels. Allowing for uncertainty aversion pushes the quarantine policy higher for longer, though it does not quite reach the true optimal response. In contrast, if the social planner initially overestimates the severity of the disease, even without uncertainty aversion the quarantine policy matches and then persists higher than the true optimal policy. With uncertainty aversion, the policy starts even higher and persists high, reflecting concerns for severe consequences even when infections and deaths never spike. Finally, when the planner's baseline model parameters match the true model, the optimal policy without ambiguity aversion is lower and less persistent than the true optimal response because the prior is somewhat diffuse. However, with ambiguity aversion the quarantine policy is persistently higher and essentially mirrors the true optimal model response. These important asymmetries in the role of uncertainty thus suggests that when there is more uncertainty about the impact of a new virus, the policymakers may want to do more to combat the spread.

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), and Hansen and Miao (2018).[4] Recent work in finance and macroeconomics has emphasized the importance of ambiguity and uncertainty, for example Hansen, Sargent, and Tallarini Jr (1999), Maenhout (2004), Garlappi, Uppal, and Wang (2007), Cogley, Colacito, Hansen, and Sargent (2008), Izhakian and Yermack (2017), Brenner and Izhakian (2018), Ai, Bansal, Guo, and Yaron (2019), and Borovička (2020). Robust control methods have also been used to study the economic impacts of climate

---

[4]There is an important and extensive theory literature on uncertainty of various forms dating to Knight (1921), Ellsberg (1961), Anscombe and Aumann (1963), and Savage (1972). Recent examples include Robson and Samuelson (2009) and Gilboa, Minardi, and Samuelson (2020).

change, as in Lemoine and Traeger (2012) and Li, Nezami Narajabad, and Temzelides (2014). In the climate change example, climate damages can have both growth effects and permanent level effects, as estimated empirically in Dell, Jones, and Olken (2012) and Colacito, Hoffmann, and Phan (2018) and analyzed theoretically in Hambel, Kraft, and Schwartz (2015) and Barnett, Brock, and Hansen (2020). 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. The paper also relates to a literature on dealing with policy uncertainty, for example Bloom (2009) and Baker, Bloom, and Davis (2016).

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 (2020), Baker, Bloom, Davis, and Terry (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 COVID19 (Manski and Molinari, 2020) and any optimal policy, whether uniform or not, will be highly sensitive to these parameters (e.g., Atkeson (2020a), Avery et al. (2020), Stock (2020)). So our quantitative results are mainly illustrative and should be interpreted with caution."*

The need to address uncertainty in models goes beyond the COVID-19 setting. Box (1976) and Cox (1995) noted the critical importance of accounting for uncertainty and unknowns in scientific analysis, stating, respectively, that *"[s]ince all models are wrong, the scientist must be alert to what is importantly wrong..."* and *"[t]he idea that complex physical, biological or sociological systems can be exactly described by a few formulae is patently absurd... construction of idealized representations that capture important stable aspects of such systems is, however, a vital part of general scientific analysis and statistical models, especially substantive ones..."* Our study addresses these issues by offering a framework for incorporating uncertainty explicitly in a wide class of macroeconomics models.

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

Figure 2: Transition Rates between States in the Augmented SIR Model.
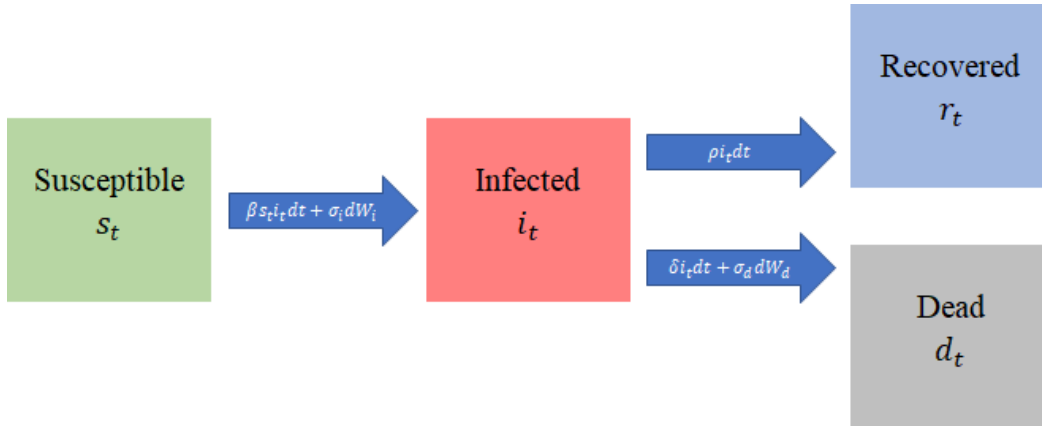
A standard SIR model is characterized by three state variables: the number susceptible individuals S_t, the number infected individuals I_t, and the number recovered individuals R_t. In addition, we include a state variable for the total population N_t, to account for deaths from the pandemic. To simplify the model solution, we use in our analysis the SIR state variables defined as fractions of the total population, i.e., $s_t = \frac{S_t}{N_t}, i_t = \frac{I_t}{N_t}, r_t = \frac{R_t}{N_t}$. 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 that increased infections has lead to increased death rates due to increased hospitalizations and therefore limited resources for treatment. So, we specify δ_t as an increasing function of i_t given by $\delta_t = \delta + \delta_+ i_t$[5]. By definition, the recovery rate of infected individuals ρ_t will depend on this specification as well.

We motivated our paper with uncertainty over these pandemic model parameters, and so departing from the standard model we introduce Brownian shocks through $W = \{W_\beta, W_\delta\}$[6], which are incorporated as parameter perturbations for β_t and δ_t with volatilities $\sigma_\beta, \sigma_\delta$, 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. Crucially, these shocks prevent the policymaker from immediately inferring the fundamental transition rates of the disease as we explore the impacts of uncertainty and ambiguity in the model.

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

$$ds_t = -\beta s_t i_t dt + s_t i_t (\delta + \delta_+ i_t) dt - s_t i_t \sigma_\beta dW_\beta + s_t i_t \sigma_\delta dW_\delta$$
$$di_t = \beta s_t i_t dt - \gamma i_t dt + i_t^2 (\delta + \delta_+ i_t) dt + \sigma_\beta s_t i_t dW_\beta - i_t \sigma_d dW_\delta + i_t^2 \sigma_\delta dW_\delta$$
$$r_t = 1 - s_t - i_t$$
$$\frac{dN_t}{N_t} = -i_t (\delta + \delta_+ i_t) dt - i_t \sigma_\delta dW_\delta$$

While we specify these state evolution equations directly here, Appendix A 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.

---

[5]This assumption and our specification is similar to the frameworks used by Alvarez et al. (2020) and Eichenbaum et al. (2020).

[6]Formally, $W \doteq \{W_t : t \geqslant 0\}$ is a multi-dimensional Brownian motion where the corresponding Brownian filtration is denoted by $\mathcal{F} \doteq \{\mathcal{F}_t : t \geqslant 0\}$ and $\mathcal{F}_t$ is generated by the Brownian motion between dates zero and t.

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:

$$ds_t = -\beta s_t i_t (1 - \zeta q_t)^2 dt + s_t i_t (\delta + \delta_+ i_t) dt - s_t i_t \sigma_\beta dW_\beta + s_t i_t \sigma_\delta dW_\delta$$
$$di_t = \beta s_t i_t (1 - \zeta q_t)^2 dt - \gamma i_t dt + i_t^2 (\delta + \delta_+ i_t) dt + \sigma_\beta s_t i_t dW_\beta - i_t \sigma_d dW_\delta + i_t^2 \sigma_\delta dW_\delta$$

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.

2.2 Economic and Public Health Model

2.2.1 Preferences, Production, and Consumption

The representative household has flow utility that depends on consumption C_t and a subjective discount rate κ, and is given by[7]

$$U_t = \kappa \log C_t$$

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. Here, A includes the capital stock, which we hold fixed. Households consume everything that is produced.

$$C_t = Y_t = AL_t$$

The labor supply is determined by the total population, which varies with shocks and

---

[7]We 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. (2020) in the appendix. Such costs, and uncertainty about these costs, should serve to enhance the results we find in our main analysis.

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:

$$L_t = N_t \left[(1 - q_t)(s_t + \phi i_t + r_t) \right]$$
$$= N_t \left[(1 - q_t)(1 - (1 - \phi)i_t) \right]$$

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.[8]

2.2.2 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 other 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 1.5 years. We provide the full details for the derivation of the model under this assumption in Appendix B, however 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.

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 uncertain planner's uncertainty-adjusted policy decisions and the difference it has on the economic and pandemic outcomes in the model. Furthermore, we will explore the differences across the scenarios of when the planner overestimates, underestimates, and correctly estimates the pandemic and the role that accounting for uncertainty has in each of the scenarios of interest

---

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

for a policymaker.

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, as well as the pandemic and economic adding up constraints. The equilibrium solution has a Markovian structure such that the value function that characterizes the solution and the optimal quarantine policy are functions of the state variables s_t, i_t, N_t. To derive the model's socially optimal outcomes, I 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 assume a discrete set Υ of possible models υ for the pandemic. Each $\upsilon \in \Upsilon$ corresponds to a set of parameters $(\beta(\upsilon), \delta(\upsilon))$. The interpretation is that each model υ comes from an existing estimate for the true pandemic model using historical data or real-time information. Each υ characterizes the state variable evolution equations as follows

$$ds_t = -\beta(\upsilon)s_t i_t(1 - \zeta q_t(\upsilon))^2 dt + s_t i_t(\delta(\upsilon) + \delta_+(\upsilon)i_t)dt - s_t i_t \sigma_\beta dW_\beta + s_t i_t \sigma_\delta dW_\delta$$

$$di_t = \beta(\upsilon)s_t i_t(1 - \zeta q_t(\upsilon))^2 dt - \gamma i_t dt + i_t^2(\delta(\upsilon) + \delta_+(\upsilon)i_t)dt + \sigma_\beta s_t i_t dW_\beta - i_t \sigma_d dW_\delta + i_t^2 \sigma_\delta dW_\delta$$

$$\frac{dN_t}{N_t} = -i_t(\delta(\upsilon) + \delta_+(\upsilon)i_t)dt - i_t \sigma_\delta dW_\delta$$

Conditional on a given υ, we derive an optimal policy $q_t(\upsilon)$ without reference to ambiguity. The standard approach then compares optimal policies across Υ. This "outside-the-model uncertainty" corresponds to the typical sensitivity analysis exercise often undertaken in the literature and does not account for the social planner making decisions under uncertainty. 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 weight the different models v, providing an uncertainty-adjusted optimal policy, based on optimization from the min-max problem.

3.1 Optimal Policy without Uncertainty

For the solutions without uncertainty, we focus on solving the social planner's problem conditional on a given model $v \in \Upsilon$. The social planner's problem is to maximize lifetime expected utility by choosing the optimal mitigation or quarantine policy $q_t(v)$. The planner's problem can be expressed as

$$V(s_t, i_t, N_t; v) = \max_{q_t(v)} E_0\Big[\int_0^T e^{-(\kappa+\lambda)t}\{\kappa \log C_t(q_t(v)) - \varphi\delta_t(v)i_t\}dt + e^{-\kappa(T-t)}\hat{V}(N_T)|v\Big]$$

subject to economic and pandemic constraints. Note that $C(q_t) = A \times L(q_t)$ and $\hat{V}(N_T)$ 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 $\hat{V}(N_T)$ and an analytical simplification for the value function given by $V(s_t, i_t, N_t; v) = \log(A) + \log N_t + v(s_t, i_t; v)$. After incorporating these simplifications, the simplified PDE we solve for the planner's problem is given by[9]

$$(\kappa + \lambda)v(s_t, i_t) = \max_{q_t} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - \varphi i_t(\delta + \delta_+ i_t) - (\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2$$
$$+ v_i\beta s_t i_t(1 - \theta q_t)^2 - v_s\beta s_t i_t(1 - \theta q_t)^2 + v_s s_t i_t(\delta + \delta_+ i_t) - v_i i_t[\gamma - (\delta + \delta_+ i_t)]i_t$$
$$+ \frac{1}{2}[v_{ss}(\sigma_d - \sigma_i)^2 s^2 i^2 + v_{ii}(\sigma_i s_t i_t - (1 - i)i_t\sigma_d)^2] - [\sigma_i^2 v_{si}]i^2 s^2$$

where we drop the v notation for brevity. The optimal choice of mitigation $q_t(v)$ is the

---

[9]See Appendix B for full details on the derivation and analytical simplification of the HJB equation.

solution to a quadratic equation resulting from the first-order condition and is given by

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

$$B = -\frac{1 + \zeta}{\zeta}$$

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

The policy function depends on the parameters associated with v, $(\beta(v), \rho(v), \delta(v), \delta_+(v))$. The standard typically stops here, focusing on the preferred model specification or perhaps includes an "outside the model" uncertainty analysis or sensitivity analysis that compares $q_t(v)$ across $v \in \Upsilon$. We include such a sensitivity or "outside the model" uncertainty analysis as part of our numerical results. However, the main takeaway we view from this comparison is the significant spread in optimal policy choices and model outcomes, only enhancing the importance of taking the uncertainty analysis further taking uncertainty "inside the model" through explicit inclusion in the planner's decision problem.

Optimal Policy with Uncertainty

Our analysis of the impact of model uncertainty is implemented through a smooth ambiguity framework that follows the decision theoretic framework developed in Hansen and Miao (2018) and applied in Barnett, Brock, and Hansen (2020). To do this, we return to the full set discrete set Υ of possible models v for the pandemic as noted above. We first specify prior probability weights for the set of models $v \in \Upsilon$, by assigning a probability weight $\pi(v)$ to each model v, satisfying

$$\pi(v) \geqslant 0 \ \ \forall v \in \Upsilon, \quad \sum_{v \in \Upsilon} \pi(v) = 1.$$

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 $\tilde{\pi}(v)$ across the set of conditional models in a way that is statistically and quantitatively reasonable by restricting the set of alternative model weights considered by the social planner based on an ambiguity aversion parameter and

the magnitude of the deviation of the distorted probability weights from the prior weights.[10] The parameter θ_a is the chosen ambiguity parameter that determines the magnitude of this penalization. Large values of θ_a imply low aversion to ambiguity, while small values of θ_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.[11]

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

$$V(s_t, i_t, N_t) = \max_{q_t} \min_{\tilde{\pi}(\upsilon)} E_0 \Big[\int_0^T e^{-(\kappa+\lambda)t} \{\kappa \log C_t(q_t) - \varphi \delta_t i_t + \sum_{\upsilon \in \Upsilon} \tilde{\pi}(\upsilon)\theta_a[\log \tilde{\pi}(\upsilon) - \log \pi(\upsilon)]\}dt + e^{-\kappa(T-t)}\tilde{V}(N_T) \Big]$$

subject to market clearing and the dynamics of the state variables relevant to the model.

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

$$(\kappa + \lambda)v(s_t, i_t) = \max_{q_t} \min_{\tilde{\pi}(\upsilon)} \kappa \log(1 - q_t) + \kappa \log(1 - (1-\phi)i_t) + \sum_{\upsilon \in \Upsilon} \tilde{\pi}(\upsilon)\theta_a[\log(\tilde{\pi}(\upsilon)) - \log(\pi(\upsilon))]$$
$$+ \sum_{\upsilon \in \Upsilon} \tilde{\pi}(\upsilon)\{-(\delta + \delta_+ i_t)i_t + \upsilon_i \beta s_t i_t (1 - \theta q_t)^2 - \upsilon_s \beta s_t i_t (1 - \theta q_t)^2$$
$$+ \upsilon_s s_t i_t (\delta + \delta_+ i_t) - \upsilon_i i_t [\gamma - (\delta + \delta_+ i_t)]\}$$
$$+ \frac{1}{2}[\upsilon_{ss}(\sigma_d - \sigma_i)^2 s^2 i^2 + \upsilon_{ii}(\sigma_i s_t i_t - (1-i)i_t\sigma_d)^2] - [\sigma_i^2 \upsilon_{si}]i^2 s^2 - \frac{1}{2}i^2\sigma_d^2$$

---

[10]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.

[11]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 significant impacts on optimal policy. In particular, we apply relative entropy penalization directly to the set of conditional value functions.

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

$$\tilde{\pi}(v) = \frac{\pi(v) \exp(-\frac{1}{\theta_a}\{v_s\mu_s(q_t) + v_i\mu_i(q_t) + \xi_N\mu_N\})}{\sum \pi(v) \exp(-\frac{1}{\theta_a}\{v_s\mu_s(q_t) + v_i\mu_i(q_t) + \xi_N\mu_N\})}$$

As the $\tilde{\pi}(v)$ 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 evolve dynamically, as was the case before. The key difference now is that the value function derivatives and drifts are what determine the distorted probabilities, and not the value function solved for each individual model. Substituting back into the objective function, our HJB after minimization can be written in a certainty-equivalent type form as follows

$$
\begin{aligned}
(\kappa + \lambda)v(s_t, i_t) = \max_{q_t} &\; \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) \\
&- \theta_a \log(\sum \tilde{\pi}(v) \exp\{-(\delta + \delta_+i_t)i_t + v_i\beta s_t i_t(1 - \theta q_t)^2 - v_s\beta s_t i_t(1 - \theta q_t)^2 \\
&+ v_s s_t i_t(\delta + \delta_+i_t) - v_i i_t[\gamma - (\delta + \delta_+i_t)]\}) \\
&+ \frac{1}{2}[v_{ss}(\sigma_d - \sigma_i)^2 s^2 i^2 + v_{ii}(\sigma_i s_t i_t - (1 - i)i_t\sigma_d)^2] - [\sigma_i^2 v_{si}]i^2 s^2 - \frac{1}{2}i^2\sigma_d^2
\end{aligned}
$$

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

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

where the terms $\hat{\beta}_t$ and $\hat{\delta}_t$ are given as before by

$$\hat{\beta}_t = \sum_v \tilde{\pi}(v)\beta(v), \quad \hat{\delta}_t = \sum_v \tilde{\pi}(v)(\delta(v) + \delta_+(v)i_t)$$

Now the planner tilts their value function and probability weights towards certain models based on state variable drifts and value function derivatives, which leads to the implied distorted model parameters which are adjusted by worst-case outcomes which the planner

uses to respond in a robustly optimal way in the face of uncertainty. Importantly, the policy function depends on the parameters associated with the various υ 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 optimal policy choice, as well as the solution for the value function.

One particular advantage to using this smooth ambiguity framework as compared to robust preferences as in Anderson, Hansen, and Sargent (2003) is that the uncertainty here is structured into alternative models as characterized by sets of key parameters. This type of structured uncertainty analysis allows us to examine 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.[12]

We note that this analysis abstracts from any form of Bayesian learning. While learning is certainly and interesting consideration to consider when thinking about the planner's optimal response to a pandemic, we find this setting valuable to consider for a number of reasons. First, 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, we view this assumption as a reasonable starting point. Second, 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.

4   Numerical Results

Here we provide numerical results from simulations based on the theoretical solutions provided above. Calibration of parameters is discussed in Appendix A.2.

---

[12]In the 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 differences in how the different types of decision theoretic frameworks impacts the structure of the planner's problem, we expect that quantitatively the results will actually be quite similar based on findings by Barnett et al. (2021) which compare such differences in a climate model uncertainty setting.

Figure 3: Outside the Model Uncertainty
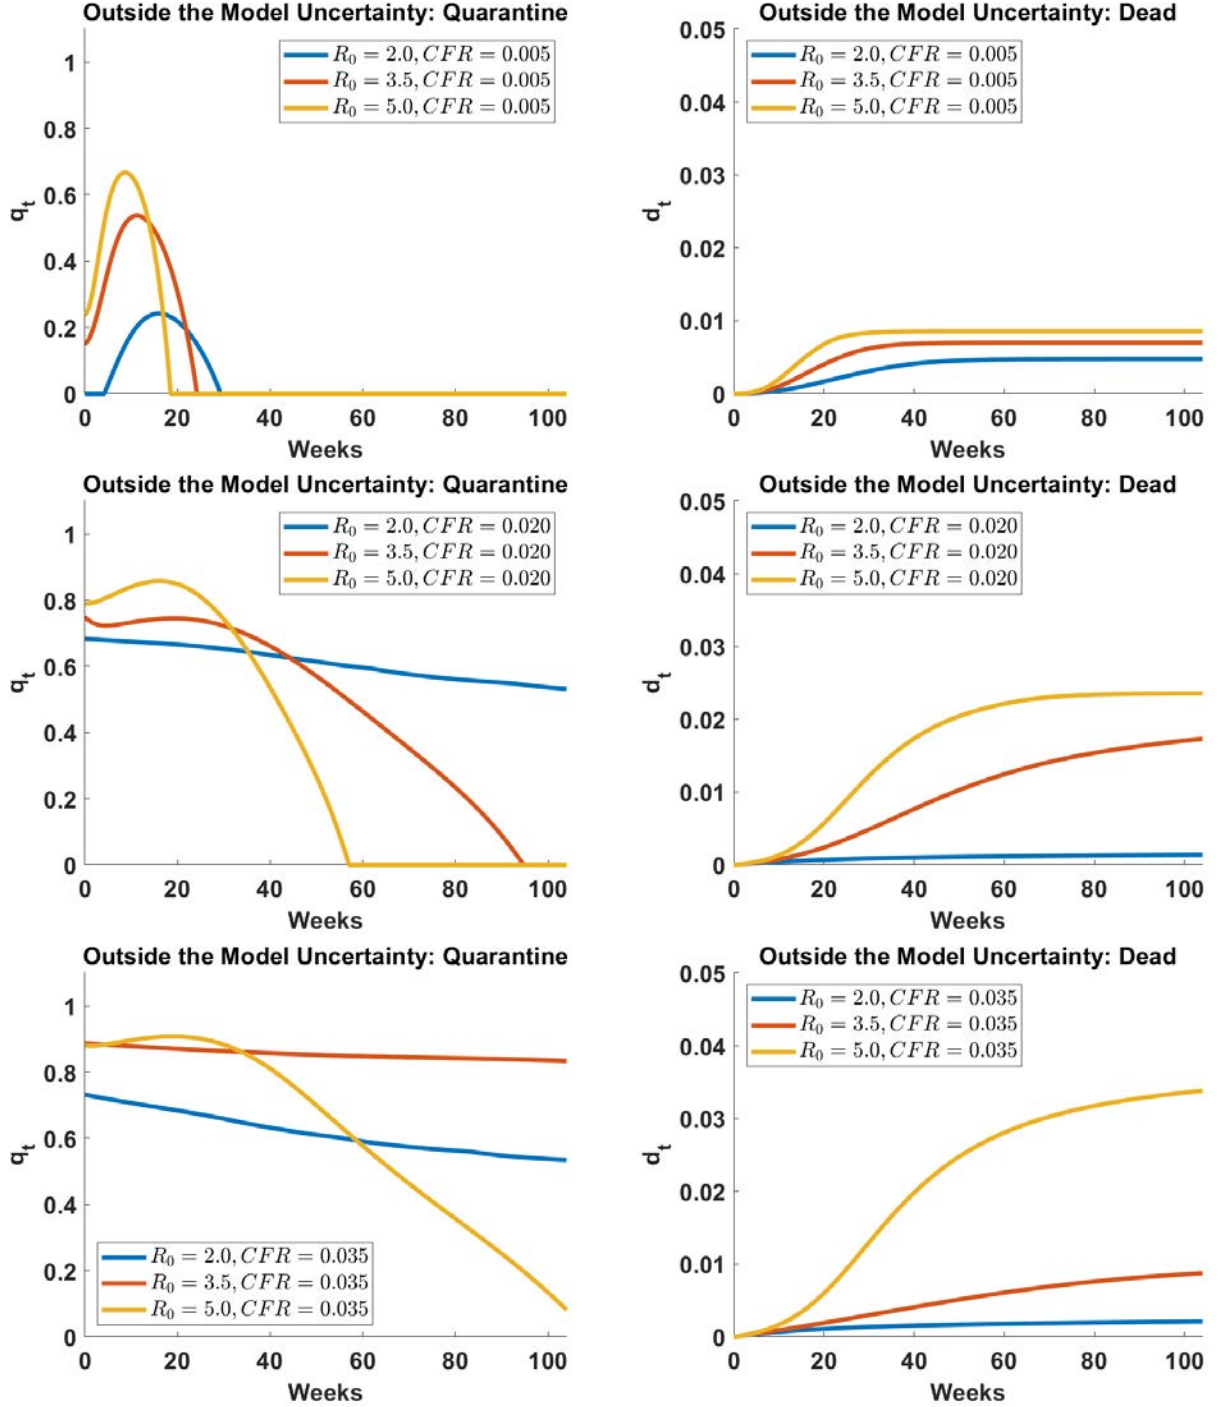
Notes: These figures show the range of possible outcomes and policy responses across nine potential models of the pandemic that vary by $\mathcal{R}_0$ and CFR. The left column shows optimal quarantine policies by model and the right column shows the fraction of the population that dies by model. The top row shows model results where $CFR = 0.005$ and $\mathcal{R}_0$ is either $2.0, 3.5$, or 5.0. The middle row shows model results where $CFR = 0.02$ and $\mathcal{R}_0$ is either $2.0, 3.5$, or 5.0. The bottom row shows model results where $CFR = 0.035$ and $\mathcal{R}_0$ is either $2.0, 3.5$, or 5.0.

16

Figure 4: Outside the Model Uncertainty
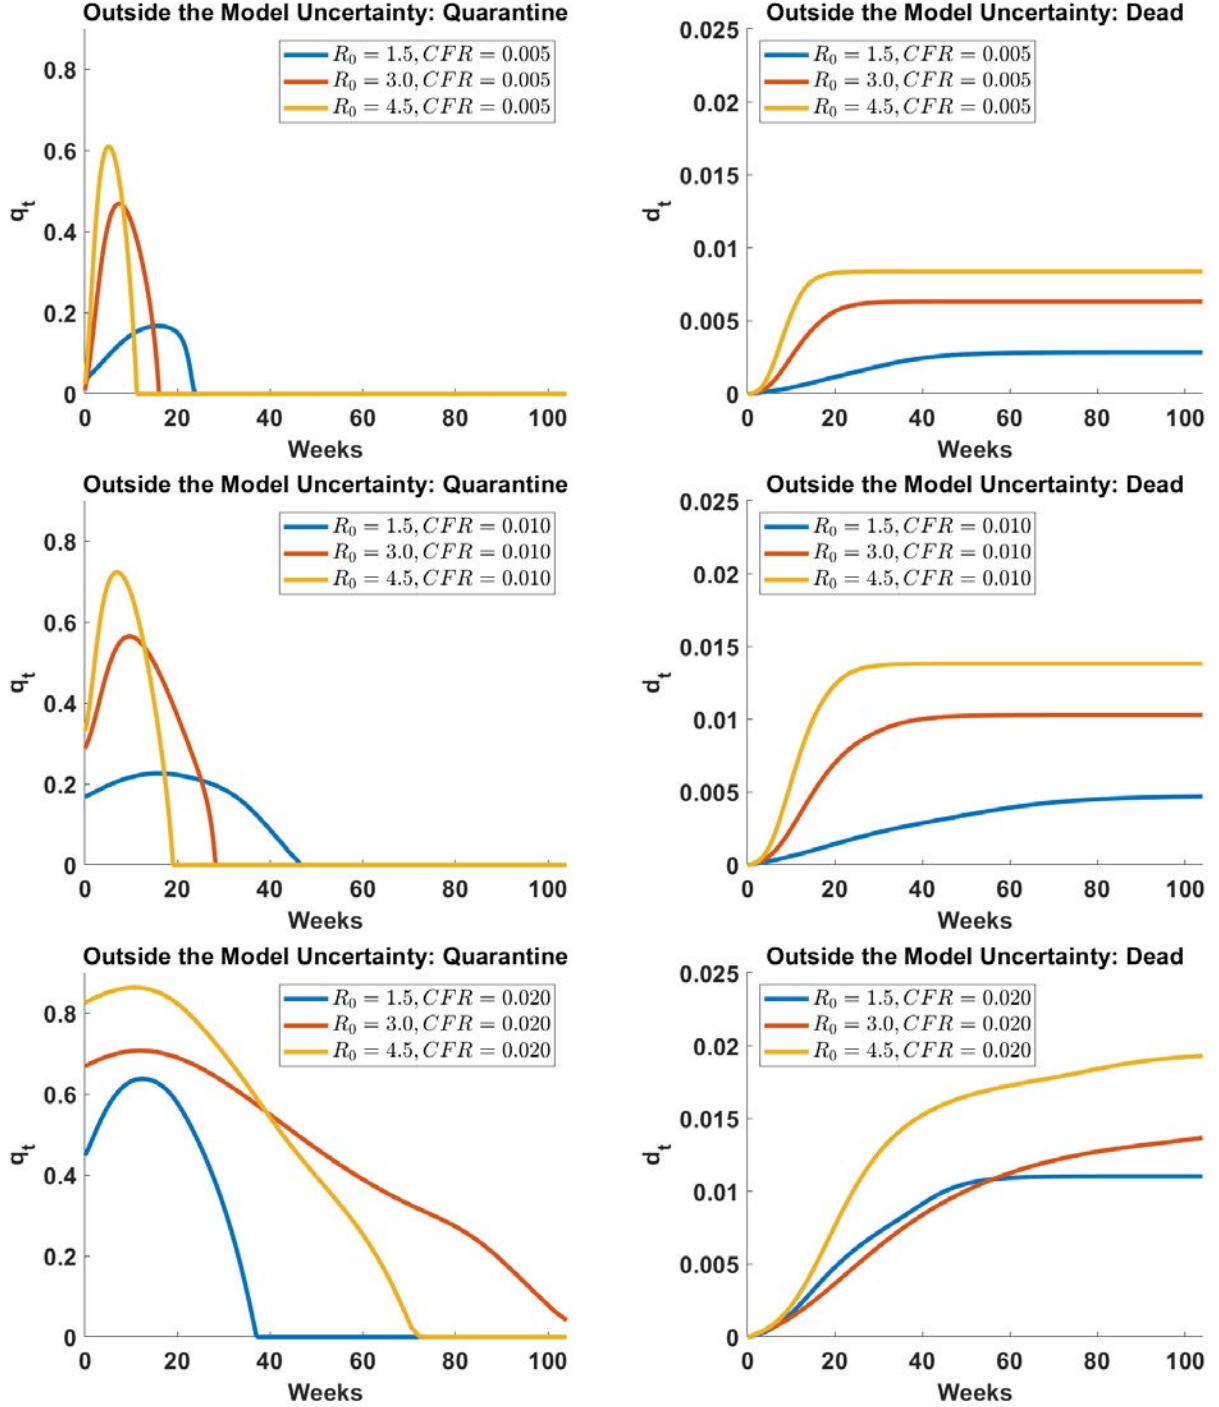
Notes: These figures show the range of possible outcomes and policy responses across nine potential models of the pandemic that vary by $\mathcal{R}_0$ and CFR. The left column shows optimal quarantine policies by model and the right column shows the fraction of the population that dies by model. The top row shows model results where $CFR = 0.005$ and $\mathcal{R}_0$ is either $2.0, 3.5$, or 5.0. The middle row shows model results where $CFR = 0.02$ and $\mathcal{R}_0$ is either $2.0, 3.5$, or 5.0. The bottom row shows model results where $CFR = 0.035$ and $\mathcal{R}_0$ is either $2.0, 3.5$, or 5.0.

17

4.1 Outside the Model Uncertainty Through 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 4 shows the spread of outcomes for d_t, and q_t across the different model cases. The spreads are across all model outcomes for $\mathcal{R}_0 \in \{2.0, 3.5, 5.0\}$ and initial CFR $\in \{0.005, 0.02, 0.035\}$.

Figure 4 indicates very different policies q and outcomes d depending on parameters. For example, the fraction of dead in the population varies by an order of magnitude, running between nearly 0 to almost 4.0%. Observe that these are death rates obtained by a policy maker who knows the true parameters and is reacting optimally, and in that sense is a best-case outcome under each scenario. However, the quarantine policies vary dramatically in magnitude and timing, suggesting 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 accounts for model uncertainty and ambiguity aversion.
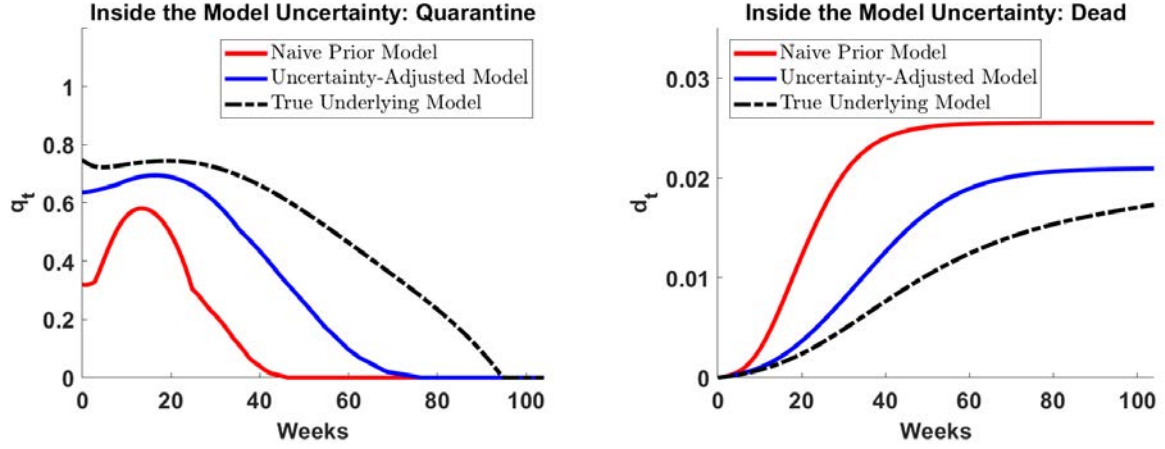
4.2 Inside the Model Uncertainty Through Smooth Ambiguity

The previous section highlighted drastically different responses and outcomes given different parameters. 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 $\mathcal{R}_0 = 2.0$ and initial CFR $= 0.02$. We then compare three sets of prior beliefs over the models: (i) underestimating the pandemic, (ii) correctly estimating the pandemic but with a disperse prior, and (iii) overestimating the pandemic. For each scenario, we compare in Figure 6 the optimal quarantine and resulting deaths from a "naive prior model" where the policymaker simply averages over optimal responses based on the assumed prior, the uncertainty-adjusted policy starting from the same prior parameters but with a somewhat diffuse prior, and finally the optimal response if the policymaker knows the true model. Additionally, Figure 7 shows how the optimal beliefs evolve over time.
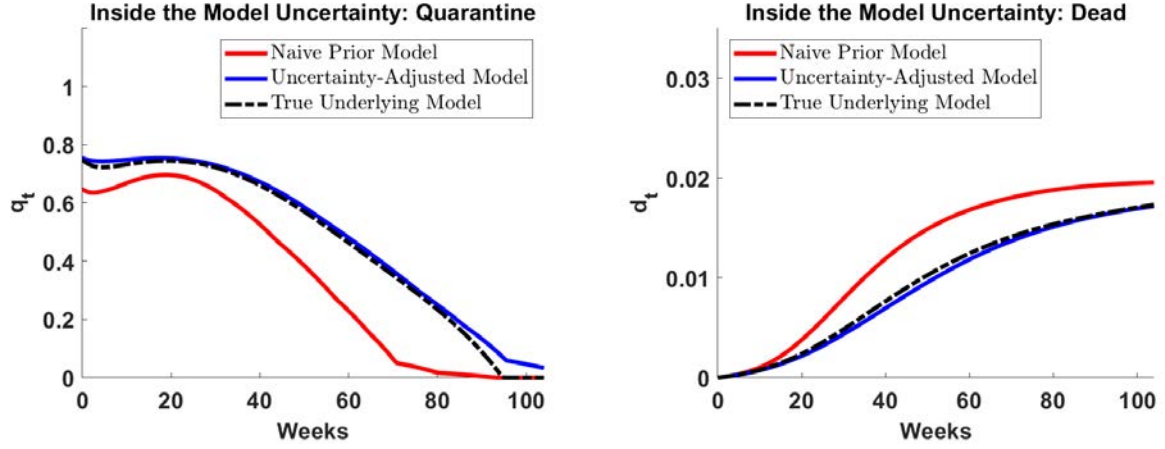
We first consider the scenario where the planner's prior underestimates the severity of the new disease by assigning a weight of 50% to the parameter set with $\mathcal{R}_0 = 2$ and initial CFR$= 0.005$ and equal weight to the rest. In this case, the naive prior model optimal response leads to much lower mitigation efforts relative to the true optimal response, as can

18

Figure 5: Optimal Mitigation Under Three Scenarios

Panel A: Underestimating the Severity of the Pandemic

Panel B: Correctly Estimating the Severity of the Pandemic

Panel C: Overestimating the Severity of the Pandemic

Notes: These figures show the fraction of the population quarantined (first column) and dead (second column) under (i) the naive prior model response (red), the uncertainty-adjusted model response (blue), and the true underlying model response (black). The three rows refer respectively to the cases where the policy maker initially under, correctly or over estimates the severity of the pandemic.

19

Figure 6: Optimal Mitigation Under Three Scenarios

Panel A: Underestimating the Severity of the Pandemic

Panel B: Correctly Estimating the Severity of the Pandemic

Panel C: Overestimating the Severity of the Pandemic

Notes: These figures show the fraction of the population quarantined (first column) and dead (second column) under (i) the naive prior model response (red), the uncertainty-adjusted model response (blue), and the true underlying model response (black). The three rows refer respectively to the cases where the policy maker initially under, correctly or over estimates the severity of the pandemic.

Figure 7: Distorted Probabilities and Death Rates Under Three Scenarios
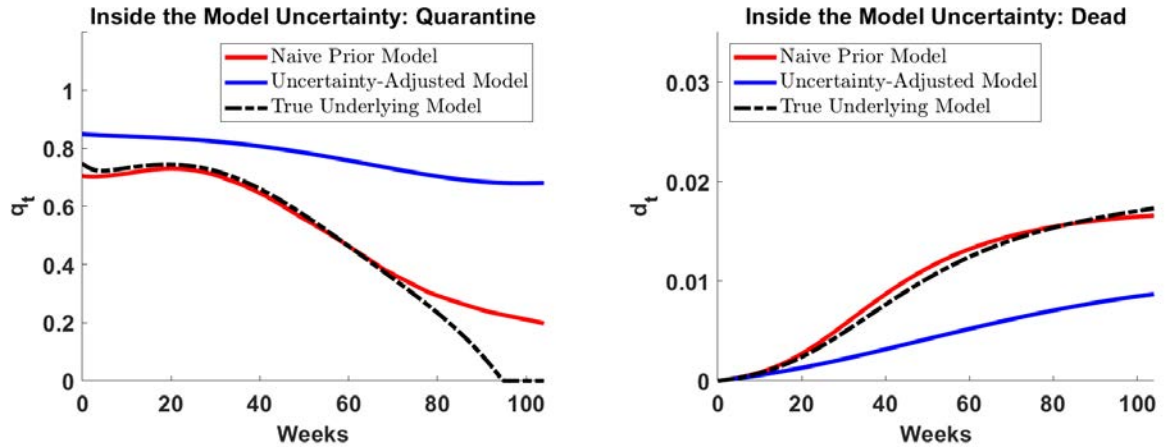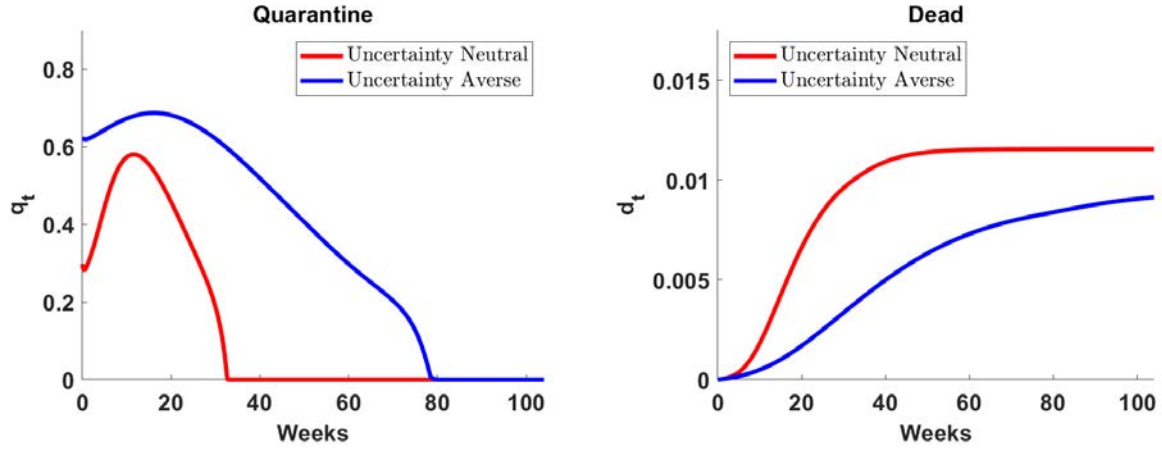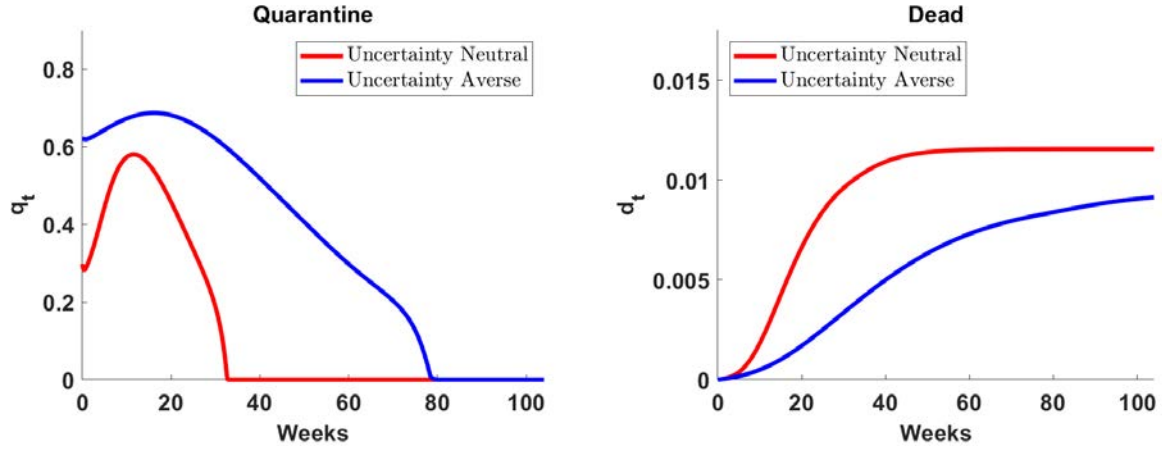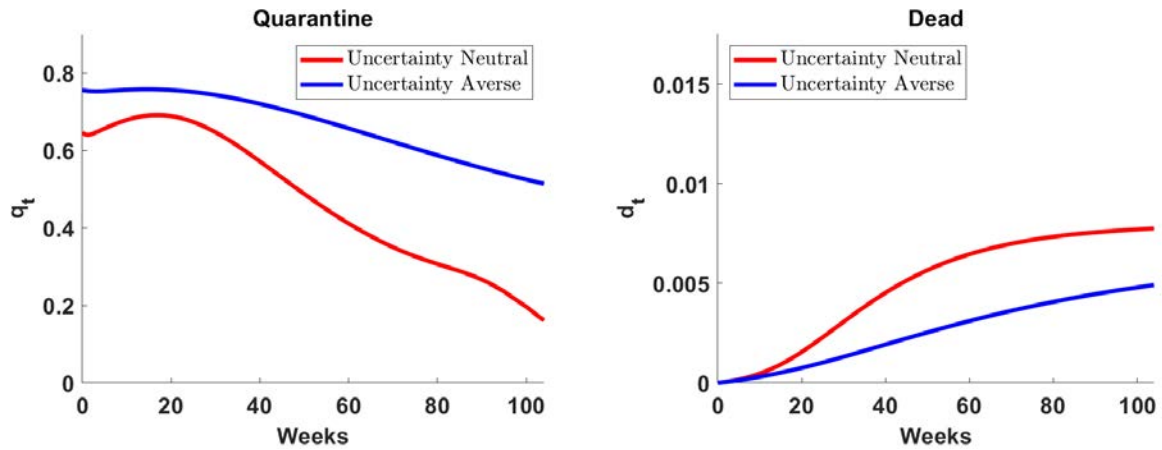
Panel A: Underestimating the Severity of the Pandemic

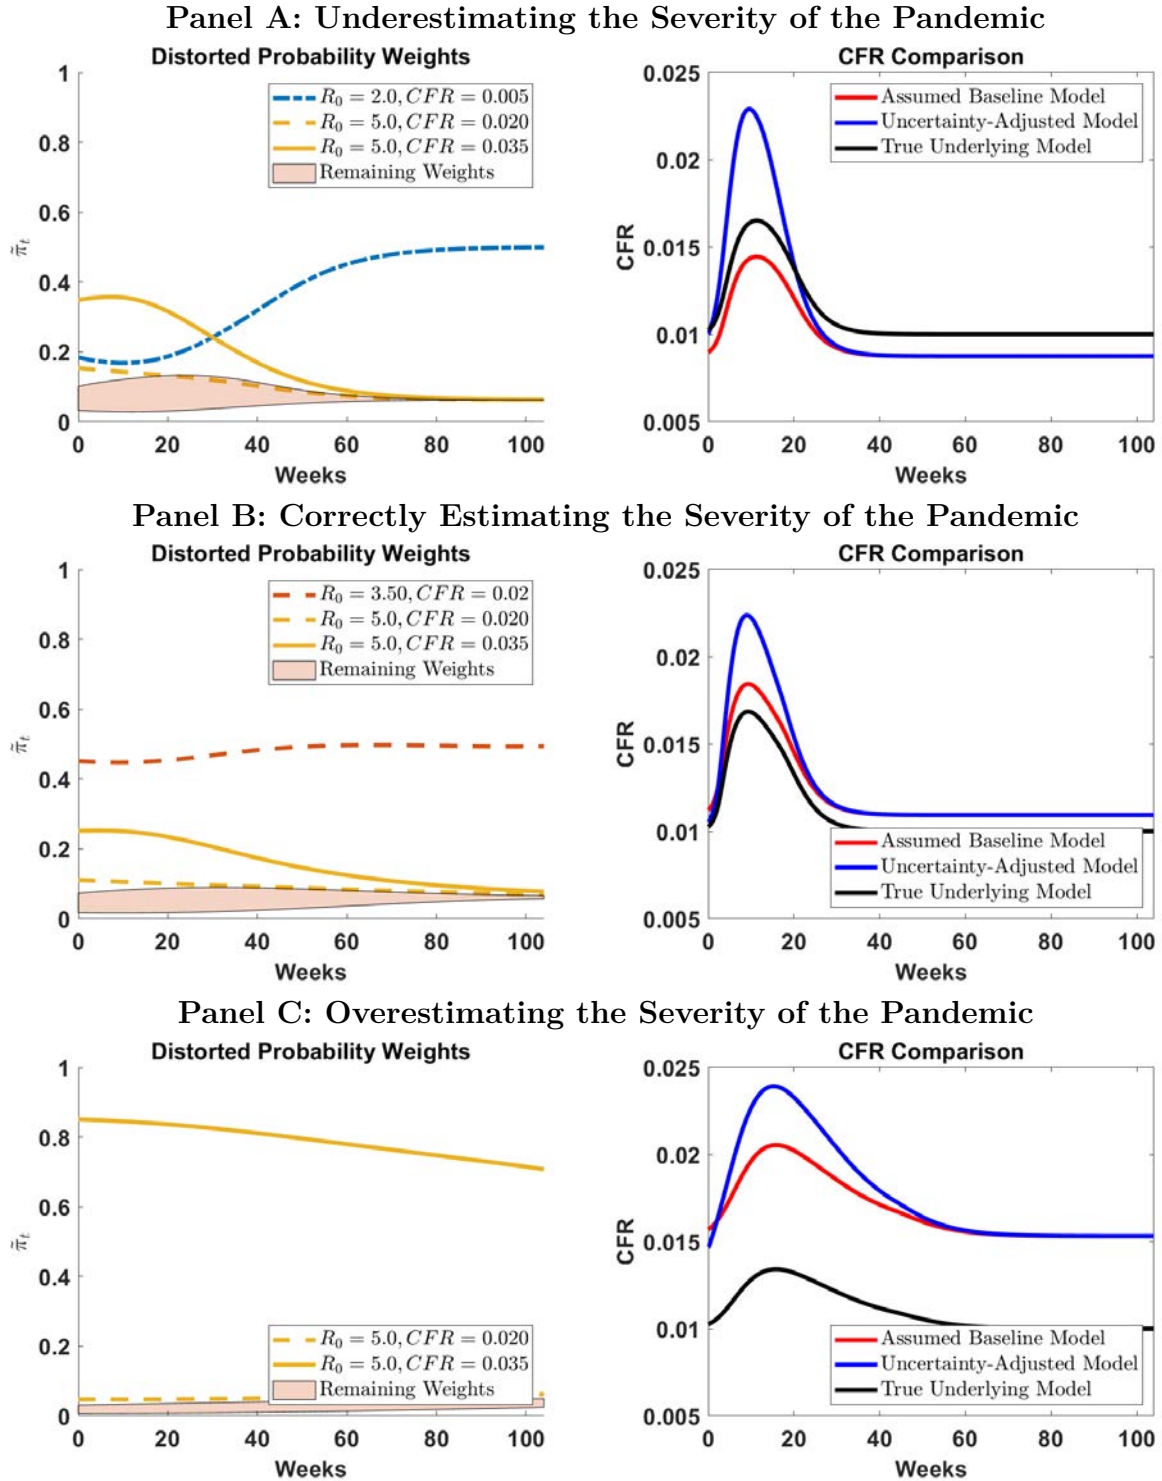## Panel B: Correctly Estimating the Severity of the Pandemic

Panel C: Overestimating the Severity of the Pandemic

Notes: The figures in the left panel show distorted probability weights $\tilde{\pi}_t$ and the figures in the right panel show the model CFR. For the distorted probability weights, the yellow solid line ($CFR = 0.035, \mathcal{R}_0 = 5.0$), and the yellow dashed line ($CFR = 0.02, \mathcal{R}_0 = 5.0$) are always included as they are two of the models that receive the highest distorted probability weight. For the top plot and middle plot, the blue dotted-dashed line ($CFR = 0.005, \mathcal{R}_0 = 2.0$) and the red dashed line ($CFR = 0.02, \mathcal{R}_0 = 3.5$) are included because they receives the highest weight for the assumed prior. The red shaded areas shows the range of distorted probability weights for the remaining models. The figures on the right panel show (i) the assumed naive prior model (red), the uncertainty-adjusted model (blue), and the true underlying model (black) values for the CFR in each case.

21

be seen in the top row of Figure 6. This results in a sharp peak in infections and a higher number of deaths relative to the optimal response. The uncertainty adjusted response brings mitigation levels close to the true optimal, yet the high level of quarantine is not sustained for as long, resulting in reduced deaths and infections compared to the naive prior model, but not to the level of the true optimal model response. Figure 7 shows how the policymaker's beliefs evolve when underestimating the pandemic. Initially, the policymaker puts a high weight on the most-severe case and remains pessimistic as infections rise. This leads to an implied CFR even higher than the true model. Eventually, as the the pandemic runs its course and then dies down, the prior weight and CFR shift toward the least-severe case that has the highest weight in the prior weighting.

We next consider the scenario where the planner correctly estimates the severity of the new disease, with a prior distribution assigning a weight of 50% to the parameter set with $\mathcal{R}_0 = 3.5$ and initial CFR= 0.02 and equal weight to the rest. The outcomes are shown in the middle row of Figure 6, with distorted beliefs shown in the middle for of Figure 7. Mitigation efforts for the naive model are higher and more persistent than when underestimating the severity of the epidemic, but still do not reach the level of the true optimal model response. The uncertainty-adjusted response, however, now mirrors the true optimal model's quarantine, and as a result, in deaths and infections as well. The probability weights start close to equal for most models, except the higher weight on the true model and an increased weight on the most-severe model shown in the solid yellow line. The impact is that the implied CFR increases for the ambiguity averse planner, persists at a higher value, and does not revert as much or as quickly as previous scenario because the prior model assumes a more severe pandemic than the first case.

In the third and final case, the planner overestimates the severity of the new disease with a prior distribution giving the model with $\mathcal{R}_0 = 5$ and initial CFR= 0.035 a weight of 50% and the remaining weight equally distributed across the other eight models. This case is shown in the bottom row of Figure 6. In this scenario, the quarantine for the naive prior model and uncertainty-adjusted model settings is high and much more persistent, incorporating uncertainty leads to excessively high quarantining. Examining the evolution of beliefs helps explain this asymmetry: when overestimating the pandemic, high weight is already placed on the worst-case model.

5 Concluding Remarks

This paper shows how to incorporate uncertainty in models of pandemics. Our main results focus on the role of uncertainty aversion in a smooth ambiguity-based decision problem, but we also show how a robust control approach would be implemented as well. With new diseases, or diseases that have only had small outbreaks, there is often significant uncertainty about key parameters which determine the overall costs of an epidemic. The results highlight important asymmetries that may be present in many robust control problems. When the assumed prior model underestimates the severity of a new threat, the planner with ambiguity aversion responds early on in a way that substantially increases the magnitude and persistence of the quarantine policy compared to response without uncertainty concerns to a level that is closer to the optimal response when the true model is known, yet the increased response does not persist as long. When the assumed prior matches the true model but with a disperse prior, the uncertainty averse planner chooses a quarantine policy that matches nearly identically the optimal true model response in terms of magnitude and persistence, even though the disbursed nature of the prior leads the non-uncertainty averse planner to under-react compared to the true optimal model response. These two cases highlight the particular importance and value of accounting for uncertainty when compared to the true optimal response. On the other hand, when the planner overestimates the pandemic severity, the overreaction is at first fairly minimal for the uncertainty-based response, though it persists at a higher level than for any other case.

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.

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.

Atkeson, A. (2020a). How deadly is covid-19? understanding the difficulties with estimation of its fatality rate. Technical report, National Bureau of Economic Research.

Atkeson, A. (2020b). What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. Technical report, National Bureau of Economic Research.

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.

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*.

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.

Brenner, M. and Y. Izhakian (2018). Asset pricing and ambiguity: Empirical evidence. *Journal of Financial Economics 130*(3), 503–531.

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.

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.

Colacito, R., B. Hoffmann, and T. Phan (2018). Temperatures and growth: A panel analysis of the united states. *Journal of Money, Credit, and Banking 51*(2-3), 2019.

Cox, D. R. (1995). Comment on "model uncertainty, data mining and statistical inference". *Journal of the Royal Statistical Society: Series A (Statistics in Society) 158*(3), 455—456.

Dell, M., B. F. Jones, and B. A. Olken (2012). Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics 4*(3), 66–95.

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.

Hambel, C., H. Kraft, and E. Schwartz (2015, March). Optimal Carbon Abatement in a Stochastic Equilibrium Model with Climate Change. Working Paper 21044, National Bureau of Economic Research.

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 (2011). Robustness and Ambiguity in Continuous Time. *Journal of Economic Theory 146*(3), 1195–1223.

Hansen, L. P., T. J. Sargent, and T. D. Tallarini Jr (1999). Robust permanent income and pricing. *Review of Economic studies*, 873–907.

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.

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 (2020). Optimal mitigation policies in a pandemic. *Working Paper*.

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*.

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*, 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.

Robson, A. J. and L. Samuelson (2009). The evolution of time preference with aggregate uncertainty. *American Economic Review 99*(5), 1925–53.

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.

Appendix A State Variable 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

$$dS_t = \phi_t N_t dt - \mu_t S_t dt - \beta_t S_t \frac{I_t}{N_t} dt$$

$$dI_t = \beta_t S_t \frac{I_t}{N_t} dt - \mu_t I_t dt - \gamma_t I_t dt$$
$$dR_t = -\mu_t R_t dt + \rho_t I_t dt$$
$$dN_t/N_t = \phi_t dt - \mu_t (S_t + I_t + R_t) dt - \delta_t I_t dt$$

where β_t is the pandemic infection rate, γ_t is the expected duration of infection, ρ_t is the pandemic recovery rate, δ_t is the pandemic death rate, μ_t is the birth rate, and ϕ_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

$$ds_t = \phi_t dt - \beta_t s_t i_t dt - s_t (\phi_t - i_t \delta_t) dt$$
$$di_t = \beta_t s_t i_t dt - (\rho_t + \delta_t) i_t dt - i_t (\phi_t - i_t \delta_t) dt$$
$$dr_t = \rho_t i_t dt - r_t (\phi_t - i_t \delta_t) dt$$

Finally to arrive at the evolution equations we use in our analysis we make two assumptions. First, we assume $\mu_t = \phi_t = 0$, 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 $\beta dt + \sigma_\beta dW_\beta$ and $\delta_t dt$ by $(\delta dt + \delta_+ i_t) + \sigma_\delta dW_\delta$. 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   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)).

We start by deriving the 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 a simple present-discounted value aggregation of linear utility given by

$$\tilde{V}(N_T) = E_0\Big[\int_T^\infty \exp(-\kappa(t-T))\kappa \log(C_t)dt|v\Big]$$

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

$$\kappa\tilde{V}(N_T) = \kappa \log A + \kappa \log N_T \iff \tilde{V}(N_T) = \log A + \log N_T$$

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

$$V(s_t, i_t, N_t; v) = \max_{q_t(v)} E_0\Big[\int_0^T \exp(-(\kappa+\lambda)t)\{\kappa \log(C_t(q_t(v)) - \varphi i_t \delta_t\}dt + \exp(-\kappa(T-t)\tilde{V}(N_T)|v]$$

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

$$
\begin{aligned}
(\kappa+\lambda)V = \max_{q_t} \ & \kappa\log(1-q_t) + \kappa\log(A) + \kappa\log(N_t) + \kappa\log(1-(1-\phi)i_t) - \varphi i_t(\delta+\delta_+ i_t) \\
& + (V_i - V_s)\beta s_t i_t(1-\theta q_t)^2 + V_s s_t i_t(\delta+\delta_+ i_t) - V_i i_t[\gamma - (\delta+\delta_+ i_t)i_t] - V_N(\delta+\delta_+ i_t)i_t N_t \\
& + \frac{1}{2}i^2 N^2 \sigma_d^2 V_{NN} + \frac{1}{2}[V_{ss}(\sigma_d-\sigma_i)^2 s^2 i^2 + V_{ii}(\sigma_i s_t i_t - (1-i)i_t\sigma_d)^2] \\
& - [\sigma_i^2 V_{si}]i^2 s^2 - V_{sN}si^2\sigma_d^2 - i^2(i-1)\sigma_d^2 V_{iN} + \lambda\log(A) + \lambda\log(N_t)
\end{aligned}
$$

Applying the analytical simplification for the value function described in the text, given by $V(s_t, i_t, N_t; v) = \log(A) + \log N_t + v(s_t, i_t; v)$, allows us to simplify the HJB equation for

the planner's problem to

$$(\kappa + \lambda)v(s_t, i_t) = \max_{q_t} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - \varphi i_t(\delta + \delta_+ i_t) - (\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2$$

$$+ v_i\beta s_t i_t(1 - \theta q_t)^2 - v_s\beta s_t i_t(1 - \theta q_t)^2 + v_s s_t i_t(\delta + \delta_+ i_t) - v_i i_t[\gamma - (\delta + \delta_+ i_t)i_t]$$

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

Note that we can alternatively use the flow utility $\kappa \log(A) + \kappa \log(N_t) + \kappa \log(1 - (1 - \phi)i_t - q_t(s_t + \phi i_t))$, where quarantine is done only on susceptible and infected workers, without any change to the algorithm that follows except for the definition of $flow$ and the expression for the optimal choice of q_t which we will see below.

We define Q as the terms involving quarantine in the HJB equation, to derive the optimal quarantine choice q_t. The equation $Q(q_t; s_t, i_t, d_t)$ is given by

$$Q(q_t; s_t, i_t d_t) = \kappa \log(1 - q_t) + (V_i - V_s)\beta s_t i_t(1 - \theta q_t)^2$$

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

$$Q' = -\kappa(1 - q_t)^{-1} - 2\theta(V_i - V_s)\beta s_t i_t(1 - \theta q_t)$$
$$Q'' = -\kappa(1 - q_t)^{-2} + 2\theta^2(V_i - V_s)\beta s_t i_t$$

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

$$q_t = \frac{-\frac{(1+\theta)}{\theta} \pm \sqrt{(\frac{(1+\theta)}{\theta})^2 - 4\{\frac{\kappa}{2\theta^2(V_i - V_s)\beta s_t i_t} + \frac{1}{\theta}\}}}{2}$$

Note that the complementary slackness condition for the first-order condition implies that $q_{max} \geqslant q_t \geqslant 0$. 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 $Q'' \geqslant 0$ 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 $Q'' \geqslant 0$ occurs the corner solution must be $q_t = 0$. Furthermore, for numerical stability we impose and upper bound constraint $q_t \leqslant 0.99$ 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

$Q'' \geqslant 0$ or $Q'' < 0$, then then solve for the new optimal q_t based on convexity or concavity, updated 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 $s_t + i_t \leqslant 1$, and therefore 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:

$$s_s \in \{0, \Delta_s, ..., 1 - \Delta_s, 1\}$$
$$i_i \in \{0, \Delta_i, ..., 1 - \Delta_i, 1\}$$

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 $s_s + i_i \leqslant 1$. 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 differencing 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 relevent 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

$$(\kappa + \lambda)v(s_t, i_t) = \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - (1 + \varphi)(\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2 - [\sigma_i^2 v_{si}]i^2 s^2$$
$$+ \frac{v(s, i + \Delta_i) - v(s, i)}{\Delta_i}\beta s_t i_t(1 - \theta q_t)^2\} - \frac{v(s, i) - v(s - \Delta_s, i)}{\Delta_s}\beta s_t i_t(1 - \theta q_t)^2$$
$$+ \frac{v(s + \Delta_s, i) - v(s, i)}{\Delta_s}s_t i_t(\delta + \delta_+ i_t) - \frac{v(s, i) - v(s, i - \Delta_i)}{\Delta_i}i_t[\gamma - (\delta + \delta_+ i_t)i_t]$$
$$+ \frac{1}{2}\frac{v(s + \Delta_s, i) + v(s - \Delta_s, i) - 2v(s, i)}{\Delta_s^2}(\sigma_d - \sigma_i)^2 s^2 i^2$$
$$+ \frac{1}{2}\frac{v(s, i + \Delta_i) + v(s, i - \Delta_i) - 2v(s, i)}{\Delta_i^2}(\sigma_i s_t i_t - (1 - i)i_t\sigma_d)^2$$

Note that we "hold constant" the cross-partial term $[\sigma_i^2 v_{si}]i^2 s^2$ by including it in the flow utility at any given iteration in the algorithm and update each these values with each iteration to find a consistent solution. For simplicity in notation after this, we define $flow \equiv \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - (1 + \varphi)(\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2 - [\sigma_i^2 v_{si}]i^2 s^2$.

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

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

Therefore, we impose

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

where the min operator is applied over the entire three-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 $s_t + i_t \leqslant 1$, 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 $s_s = 0, 0 < i_i < i_{max}$:

$$v(s_t, i_t) = \max_{q_t} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - (1 + \varphi)(\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2$$

$$+ (1 - (\kappa + \lambda)dt)[v(s, i - \Delta_i) - v(s, i)]\{\frac{i_t[\gamma - (\delta + \delta_+ i_t)i_t]}{(1 - (\kappa + \nu)dt)}\frac{dt}{\Delta_i}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s, i + \Delta_i) - 2v(s, i) + v(s, i - \Delta_i)]\{\frac{\frac{1}{2}(-(1-i)i_t\sigma_d)^2}{(1 - (\kappa + \lambda)dt)}\frac{dt}{\Delta i^2}\}$$

$$+ (1 - (\kappa + \lambda)dt)v(s, i)$$

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

$$v(s_t, i_t) = \max_{q_t} \kappa \log(1 - (1 - \phi)) - (1 + \varphi)(\delta + \delta_+) - \frac{1}{2}\sigma_d^2$$

$$+ (1 - (\kappa + \nu)dt)[v(s, i - \Delta_i) - v(s, i)]\{\frac{\gamma - (\delta + \delta_+)}{(1 - (\kappa + \lambda)dt)}\frac{dt}{\Delta_i}\}$$

$$+ (1 - (\kappa + \lambda)dt)v(s, i)$$

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

$$v(s_t, i_t) = \max_{q_t} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - (1 + \varphi)(\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2 - [\sigma_i^2 v_{si}]i^2 s^2$$

$$+ (1 - (\kappa + \lambda)dt)[v(s - \Delta_s, i) - v(s, i)]\{\frac{\beta s_t i_t(1 - \theta q_t)^2}{(1 - (\kappa + \lambda)dt)}\frac{dt}{\Delta_s}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s + \Delta_s, i) - v(s, i)]\{\frac{s_t i_t(\delta + \delta_+ i_t)}{(1 - (\kappa + \lambda)dt)}\frac{dt}{\Delta_s}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s, i - \Delta_i) - v(s, i)]\{\frac{i_t[\gamma - (\delta + \delta_+ i_t)i_t]}{(1 - (\kappa + \lambda)dt)}\frac{dt}{\Delta_i}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s, i + \Delta_i) - v(s, i)]\{\frac{\beta s_t i_t(1 - \theta q_t)^2}{(1 - (\kappa + \lambda)dt)}\frac{dt}{\Delta_i}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s + \Delta_s, i) - 2v(s, i) + v(s - \Delta_s, i)]\{\frac{\frac{1}{2}(\sigma_d - \sigma_i)^2 s^2 i^2}{(1 - (\kappa + \lambda)dt)}\frac{dt}{\Delta_s^2}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s, i + \Delta_i) - 2v(s, i) + v(s, i - \Delta_i)]\{\frac{\frac{1}{2}(\sigma_i s_t i_t - (1 - i)i_t \sigma_d)^2}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta i^2}\}$$

$$+ (1 - (\kappa + \lambda)dt)v(s, i)$$

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

$$v_i - v_s \approx \frac{v(s, i + \Delta_i, d) - v(s, i, d)}{\Delta_i} - \frac{v(s, i, d) - v(s - \Delta_s, i, d)}{\Delta_s}$$

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

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

$$v(s_t, i_t) = \max_{q_t} \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - (1 + \varphi)(\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2 - [\sigma_i^2 v_{si}]i^2 s^2$$

$$+ (1 - (\kappa + \lambda)dt)[v(s - \Delta_s, i + \Delta_i) - v(s, i)]\{\frac{\beta s_t i_t (1 - \theta q_t)^2}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_s}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s + \Delta_s, i - \Delta_i) - v(s, i)]\{\frac{s_t i_t (\delta + \delta_+ i_t)}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_s}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s, i - \Delta_i) - v(s, i)]\{\frac{i_t[\gamma - (\delta + \delta_+ i_t)i_t]}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_i}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s + \Delta_s, i - \Delta_i) - 2v(s, i) + v(s - \Delta_s, i)]\{\frac{\frac{1}{2}(\sigma_d - \sigma_i)^2 s^2 i^2}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta_s^2}\}$$

$$+ (1 - (\kappa + \lambda)dt)[v(s - \Delta_s, i + \Delta_i) - 2v(s, i) + v(s, i - \Delta_i)]\{\frac{\frac{1}{2}(\sigma_i s_t i_t - (1 - i)i_t \sigma_d)^2}{(1 - (\kappa + \lambda)dt)} \frac{dt}{\Delta i^2}\}$$

$$+ (1 - (\kappa + \lambda)dt)v(s, i)$$

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

$$v_i - v_s \approx \frac{v(s - \Delta_s, i + \Delta_i) - v(s, i)}{\Delta_s}$$

$$v_{si} \approx \frac{v(s, i) - v(s, i - \Delta_i) - v(s - \Delta_s, i) + v(s - \Delta_s, i - \Delta_i)}{\Delta_s \Delta_i}$$

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 is in line with Alvarez et al. (2020) and consistent with some recent research that exposure to related coronaviruses may provide some people with at least partial immunity[13], and insures stability of the simulations as well. The initial values we use are given by $s_0 = 1 - \epsilon_1 - \epsilon_2$, $i_0 = \epsilon_1 = 0.01$, $r_0 = 0.02$, and $d_0 = 0$.

Appendix C   Parameter Values

This appendix 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 to one and set productivity to $A = 20/12 = 1.667$ so that output in the non-pandemic version of the model $(A \times L)$ 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 $\kappa = 0.0075$ for the baseline analysis. We assume the expected arrival time for a vaccine is one year, so that the value of the arrival rate is given by $\lambda = 1 \times \frac{1}{12} = 0.056$.

For the pandemic model parameters, we use values from various studies (including Korolev (2020), Atkeson (2020b), Atkeson (2020a), 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 the birth rate $\mathcal{R}_0$, 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 one month. The set of underlying models used in our analysis use initial values of CFR in the set $\{0.005, 0.02, 0.035\}$ and values of $\mathcal{R}_0$ in the set $\{2.0, 3.5, 5.0\}$. 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

---

[13]See articles in Science, Nature, and medRxiv for examples.

range of values across these different studies. The value of ϕ is consistent with estimated values for the fraction of infected individuals who are asymptomatic given by the CDC COVD-19 Pandemic Planning Scenarios website. The value for quarantine effectiveness, ζ, is chosen to match Alvarez, Argente, and Lippi (2020). For the volatilities σ_i and σ_d, we use data from the Center for Systems Science and Engineering in the Whiting School of Engineering at Johns Hopkins University[14] to calculate empirical counterparts for these.

Finally, we must also specify values for the uncertainty parameter in our model θ_a for smooth ambiguity. Our main value main of the uncertainty parameter, $\theta_a = 0.0275$, imposes a significant amount of uncertainty aversion to demonstrate the potential magnitude of uncertainty impacts. This value can be difficult to interpret on its own, and are best interpreted by way of the conditional relative entropy values implied by these parameter choices and the distorted model probabilities provided in the analysis to determine if they are reasonable. We do not provide those values here, but they can be provided upon request from the authors. Furthermore, anecdotal evidence on model spreads implied by the recent estimates of COVID-19 parameter values which guide the values we use ensure that our distorted values remain within a reasonable region.

Table A.2: Parameter Values

Parameters	Variable	Value
Subjective Discount Rate	κ	0.0075
Non-Pandemic Output	$A \times \bar{L}$	1.667
Infection Severity	ϕ	0.4
Quarantine Effectiveness	ζ	0.5
Arrival Rate of Vaccine	λ	0.056
Reproduction Number	$\mathcal{R}_0$	$\{2.0, 3.5, 5.0\}$
Initial Case Fatality Rate	CFR	$\{0.005, 0.02, 0.035\}$
Infection Half Life	γ	$30/18$
Death Rate Convexity	δ_+	$5 \times \frac{CFR}{\gamma}$
Volatility	σ_i, σ_d	$\{0.029, 0.014\}$
Ambiguity Parameter	θ_a	$\{0.0015, \infty\}$

---

[14]This data is available through the CSSE GitHub repo.

Appendix D   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 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 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. (2020):

$$\chi(\iota) \;\; = \varphi(\delta + \delta_+ \iota)\iota$$

The value of φ represents the valuation the planner places on the deaths from the pandemic and $(\delta + \delta_+ i_t)i_t$ 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

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

where χ_i and χ_{ii} are the first and second derivatives, respectively of $\chi(i_t)$ and $\mu_i(s_t, i_t; q_t, \upsilon)$ and $\sigma_i(s_t, i_t; q_t, \upsilon)$ 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 nonpecuniary 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 uncertainty and its implications for the choice of optimal quarantine policy.

Productivity Costs of Mitigation and Uncertainty

An important extension to consider is the possibility of additional costs of quarantine measures on output, beyond the consumption-equivalent costs that result from reduced labor. In addition, a potential consequence of mitigation efforts is that it could lead to reduction in productivity. In particular, Barrot, Grassi, and Sauvagnat (2020) note that social distancing measures could lead to a reduction in GDP growth. We model this formally by extending our expression of productivity to be

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

$$dz_t = \mu_z dt - \frac{1}{2}\sigma_z^2 dt + \sigma_z dW_t$$

This extends productivity to follow standard geometric Brownian motion growth as is commonly used throughout economics and financial modeling. However, we add to this a term to account for reduced growth resulting from social distancing quarantine measures, which we will calibrate to fit the estimates provided by Barrot, Grassi, and Sauvagnat (2020). This additional term augments the process for z to now be

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

$$dz_t = \mu_z dt - \tilde{a}q_t^{\tilde{b}}dt - \frac{1}{2}\sigma_z^2 dt + \sigma_z d\hat{W}_t$$

The cost provides an additional impact from quarantine reflected in not only level impacts but also growth implications for quarantine measures. In addition, as there exists substantial uncertainty about the long-term economic consequences of "shutting down the economy" in this manner, we can allow for this additional channel of model uncertainty as we have done with the pandemic model, allowing for alternative values of $\tilde{a}, \tilde{b}$ to be specified and part of our υ conditional models so that dz_t is given b

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

$$dz_t = \mu_z dt - \tilde{a}(\upsilon)q_t^{\tilde{b}(\upsilon)}dt - \frac{1}{2}\sigma_z^2 dt + \sigma_z dW_t$$

This additional channel of uncertainty will interact with our existing uncertainties and will potentially have meaningful implications for the social planner's optimal policy response.

Uncertainty Through Robustness

Though our analysis mainly 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 literature[15]. 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:

$$
\begin{aligned}
ds_t &= -(\beta + \sigma_\beta h_t)s_t i_t dt + s_t i_t(\delta + \delta_+ i_t + \sigma_\delta h_t)dt - s_t i_t \sigma_\beta dW_\beta + s_t i_t \sigma_\delta dW_\delta \\
di_t &= (\beta + \sigma_\beta h_t)s_t i_t dt - \gamma i_t dt + i_t^2(\delta + \delta_+ i_t + \sigma_\delta h_t)dt + \sigma_\beta s_t i_t dW_\beta - i_t(1 - i_t)\sigma_d dW_\delta \\
r_t &= 1 - s_t - i_t \\
dN_t &= -i_t N_t(\delta + \delta_+ i_t + \sigma_\delta h_t)dt - i_t N_t \sigma_\delta dW_\delta
\end{aligned}
$$

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,

---

[15]Detailed 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).

alternative models under consideration by the social planner are of the form

$$
\begin{aligned}
ds_t &= -(\beta + \sigma_\beta h_t)s_t i_t (1 - \zeta q_t)^2 dt + s_t i_t (\delta + \delta_+ i_t + \sigma_\delta h_t) dt - s_t i_t \sigma_\beta dW_\beta + s_t i_t \sigma_\delta dW_\delta \\
di_t &= (\beta + \sigma_\beta h_t)s_t i_t (1 - \zeta q_t)^2 dt - \gamma i_t dt + i_t^2 (\delta + \delta_+ i_t + \sigma_\delta h_t) dt + \sigma_\beta s_t i_t dW_\beta - i_t (1 - i_t) \sigma_d dW_\delta \\
r_t &= 1 - s_t - i_t \\
dN_t &= -i_t N_t (\delta + \delta_+ i_t + \sigma_\delta h_t) dt - i_t N_t \sigma_\delta dW_\delta
\end{aligned}
$$

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.

We can interpret the drift perturbations for misspecification directly as parameter misspecifications and altered model parameters of the form

$$
\begin{aligned}
ds_t &= -\tilde{\beta}_t s_t i_t (1 - \zeta q_t)^2 dt + s_t i_t \tilde{\delta}_t dt - s_t i_t \sigma_\beta dW_\beta + s_t i_t \sigma_\delta dW_\delta \\
di_t &= \tilde{\beta}_t s_t i_t (1 - \zeta q_t)^2 dt - \gamma i_t dt + i_t^2 \tilde{\delta}_t dt + \sigma_\beta s_t i_t dW_\beta - i_t (1 - i_t) \sigma_d dW_\delta \\
dN_t &= -i_t N_t \tilde{\delta}_t dt - i_t N_t \sigma_\delta dW_\delta
\end{aligned}
$$

where $\tilde{\beta}_t = \beta + h_t \sigma_\beta$ and $\tilde{\delta}_t = \delta + \delta_+ i_t + h_t \sigma_\delta$. 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 will 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 θ_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 θ_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

$$V(s_t, i_t, N_t) = \max_{q_t} \min_{h_t} E_0\Big[\int_0^T e^{-(\kappa+\lambda)t}\{\kappa \log(C_t(q_t) - \varphi\delta_t i_t + \frac{\theta_m}{2}|h_t|^2\}dt + e^{-\kappa(T-t)}\tilde{V}(N_T)|\upsilon\Big]$$

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

$$
\begin{aligned}
(\kappa + \lambda)v(s_t, i_t) = \max_{q_t} \ & \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - (1 + \varphi)(\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2 + \frac{\theta_m}{2}|h_t|^2 \\
& + v_i\beta s_t i_t(1 - \theta q_t)^2 - v_s\beta s_t i_t(1 - \theta q_t)^2 + v_s s_t i_t(\delta + \delta_+ i_t) - v_i i_t[\gamma - (\delta + \delta_+ i_t)] \\
& + \left[-s_t i_t\sigma_\beta v_s + s_t i_t\sigma_\delta v_s + \sigma_\beta s_t i_t v_i - i_t(1 - i_t)\sigma_d v_i - i_t\sigma_\delta\right]h_t \\
& + \frac{1}{2}\left[v_{ss}(\sigma_d - \sigma_i)^2 s^2 i^2 + v_{ii}(\sigma_i s_t i_t - (1 - i)i_t\sigma_d)^2\right] - [\sigma_i^2 v_{si}]i^2 s^2
\end{aligned}
$$

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

$$|h_t|^2 = \frac{1}{\theta_m^2}\big[(v_i - v_s)^2(s_t i_t\sigma_\beta)^2 + (s_t v_s - (1 - i_t)v_i - 1)^2(i_t\sigma_\delta)^2\big]$$

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

$$
\begin{aligned}
(\kappa + \lambda)v(s_t, i_t) = \max_{q_t} \ & \kappa \log(1 - q_t) + \kappa \log(1 - (1 - \phi)i_t) - (1 + \varphi)(\delta + \delta_+ i_t)i_t - \frac{1}{2}i^2\sigma_d^2 \\
& - \frac{1}{2\theta_m}\big[(v_i - v_s)^2(s_t i_t\sigma_\beta)^2 + (s_t v_s - (1 - i_t)v_i - 1)^2(i_t\sigma_\delta)^2\big] \\
& + v_i\beta s_t i_t(1 - \theta q_t)^2 - v_s\beta s_t i_t(1 - \theta q_t)^2 + v_s s_t i_t(\delta + \delta_+ i_t) - v_i i_t[\gamma - (\delta + \delta_+ i_t)]
\end{aligned}
$$

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

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 optimal mitigation policy takes the same functional form as before. 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, and dead populations, represented by V_s, V_i, V_d. Though we do not report results from this approach, they can easily be solved for numerically and should provide similar qualitative findings to those under the smooth ambiguity approach.

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 $q_t(\upsilon)$ and value function $v(s_t, i_t; \upsilon)$ for each conditional model, we first specify a prior distribution to the set of models $\upsilon \in \Upsilon$, by assigning a probability weight $\pi(\upsilon)$ to each model υ, satisfying

$$\pi(\upsilon) \geqslant 0 \quad \forall \upsilon \in \Upsilon, \quad \sum_{\upsilon \in \Upsilon} \pi(\upsilon) = 1.$$

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 $\tilde{\pi}(\upsilon)$ 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.[16] The parameter θ_a is chosen to determine

---

[16]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.

the magnitude of this penalization. Large values of θ_a imply low aversion to ambiguity, while small values of θ_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.[17]

This new, second-stage problem for the planner is a minimization problem, where the minimization is made over possible distorted probability weights $\tilde{\pi}(v)$ which are constrained by θ_a based on the solutions to the v 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.[18] 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 that are robust to alternative models. The second-stage minimization problem is given by the solution to the following problem

$$\tilde{V}_t = \min_{\tilde{\pi}(v)} \sum_{v \in \Upsilon} \tilde{\pi}(v)(V(v) + \theta_a[\log(\tilde{\pi}(v)) - \log(\pi(v))])$$

$$\text{subject to} \sum_{v \in \Upsilon} \pi(v) = \sum_{v \in \Upsilon} \tilde{\pi}(v) = 1$$

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

$$\tilde{\pi}(v) = \pi(v) \frac{\exp(-\frac{1}{\theta_a} V(v))}{\sum \pi(v) \exp(-\frac{1}{\theta_a} V(v))}$$

As the $\tilde{\pi}(v)$ in the model are optimally determined and state dependent, the magnitude of the ambiguity adjustment considered by the social planner when making optimal policy

---

[17]See 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.

[18]One 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.

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 υ 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 υ conditional optimal mitigation policies $q_t(\upsilon)$ using the distorted probability weights. The magnitude of the weight given to each υ conditional model informs the planner on how to use the υ 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

$$\tilde{q}_t = \sum_{\upsilon \in \Upsilon} \tilde{\pi}(\upsilon) q_t(\upsilon), \quad \hat{\beta}_t = \sum_{\upsilon \in \Upsilon} \tilde{\pi}(\upsilon) \beta(\upsilon), \quad \hat{\delta}_t = \sum_{\upsilon \in \Upsilon} \tilde{\pi}(\upsilon)(\delta(\upsilon) + \delta_+(\upsilon))$$

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.