PNAS
www.pnas.org

Supplementary Information for

Epidemic Responses Under Uncertainty

Michael Barnett, Greg Buchak and Constantine Yannelis

Corresponding Author: Michael Barnett.
E-mail: michael.d.barnett@asu.edu

This PDF file includes:

Supporting Information Text

1. SIR Model Dynamics and Volatility

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

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

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

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

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

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

2. Detection Error Probabilities

To quantify the magnitude of uncertainty allowed for in our analysis based on the choice of ξ_a, we use detection error probabilities as proposed by (1) and (2). These probabilities, now commonly used in the literature on model uncertainty, are based on the model discrimination bounds proposed by (3) and (4). The bound quantifies the probability of making a type I or type II error when choosing between two models. In our setting, the two models are the baseline model where state dynamics depend on the prior probabilities and the worst-case model where state dynamics depend on the distorted posterior.

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

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

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

and under the distorted prior, the evolution is given by

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

Note that $\boldsymbol{\sigma}$ is the matrix of volatilities for each state and shock and $\boldsymbol{\mu}$ and $\tilde{\boldsymbol{\mu}}$ are the vector of drifts for each state variable weighted by the prior weights $\pi(\theta)$ and the uncertainty adjusted weights $\tilde{\pi}(\theta)$ across each model θ, respectively. Following (5), the drift distortion implied by the structured ambiguity across model parameters is given by

$$\boldsymbol{h_t} = (\boldsymbol{\sigma}'\boldsymbol{\sigma})^{-1} \boldsymbol{\sigma}' (\tilde{\boldsymbol{\mu}} - \boldsymbol{\mu})$$

Given the drift distortions, we solve for the detection error probability through simulation methods as in (2), (1), and (6). This provides us with a state-dependent, conditional model discrimination measure to determine whether our choice of ξ_a is reasonable in that the detection error probability would be small enough that we would be confident in a statistically significant sense of which model is the true model (i.e., $< 10\%$). To calculate the detection error probabilities, we use discrete-time approximations of the state variables of interest. Define $x_t \equiv [s_t, i_t, N_t, z_t]$ and $\tilde{x}_t \equiv \left[\tilde{s}_t, \tilde{i}_t, \tilde{N}_t, \tilde{z}_t\right]$, where x_t evolves under the baseline prior and $\tilde{x}_t$ evolves under the distorted distribution as follows:

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

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

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

Note that we set $T = 12$ months as mentioned previously, and we simulate 100,000 pathways for $\boldsymbol{x}_{t+1}$ and $\tilde{\boldsymbol{x}}_{t+1}$ at each point in time of our simulated model solution results. Therefore, the detection error probability calculated at each point in time of our simulated model results is based on information from 100,000 independent and uniquely calculated pathways with equal sample length. The value of 1/2 in step 5 comes from choosing an equal-weighted prior on each type of error p_A and p_B, where p_A represents the probability of choosing the distorted model as the model that generated the sample path when it is actually the baseline model and p_B represents the probability of choosing the baseline model as the model that generated the sample path when it is actually the distorted model. Thus, a detection error probability of 0.5 means that there is essentially no way to statistically discriminate between the distorted model and the baseline model, while a detection error probability close to 0.0 means that one can statistically discriminate between the distorted model and the baseline model with near certainty.

3. Model Extensions

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

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

The nonpecuniary losses for deaths from the pandemic x_t account for losses beyond the economic costs of a reduced labor force, and therefore reduced final output production, from quarantine measures and death. Because these costs are determined by deaths from the pandemic, there is a mapping between infections and nonpecuniary costs of the form $x_t = \chi(i_t)$, where we assume a functional form similar to (7), (8), and (9):

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

where $\delta_t = \delta + \delta_+\iota$ as noted in the main text. The value of φ represents the valuation the planner places on the deaths from the pandemic and $\delta_t i_t$ 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

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

where χ_i and χ_{ii} are the first and second derivatives, respectively of $\chi(i_t)$ and $\mu_i(s_t, i_t; q_t, \theta)$ and $\sigma_i(s_t, i_t)$ are the drift and volatility, respectively, of the state variable for infections i_t. Note optimal quarantine policy influences the drift of i_t which directly enters the evolution of x_t. Therefore, the optimal choice for q_t will include an explicit adjustment for these 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 to uncertainty and its implications for the choice of optimal quarantine policy.

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

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

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

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

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

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

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

For the uncertainty analysis to be reasonable, we restrict the set of alternative models considered by the social planner to those that are difficult to distinguish from the baseline model using statistical methods and past data. A penalization term based on the conditional relative entropy measure of model distance is used to accomplish this. The parameter ξ_m is chosen to determine the magnitude of this penalization. We have defined relative entropy previously, and note that (11) provides complete details about relative entropy use in a robust preferences setting. Again, relative entropy means we are only considering relatively small, though potentially significant, distortions from the baseline model.

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

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

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

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

*Detailed explanations of robust preference problems and axiomatic treatment of such formulations using penalization methods are given by (10), (2), (11), (12), (13) and (14).

which characterizes the socially optimal solution is now given by

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

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

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

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

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

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

Uncertainty Through Alternative Smooth Ambiguity Formulation. Here we incorporate uncertainty using the decision theoretic framework developed in (14). Given an optimal policy $q_t(\theta)$ and value function $v(s_t, i_t; \theta)$ for each conditional model, we first specify a prior distribution to the set of models $\theta \in \Theta$, by assigning a probability weight $\pi(\theta)$ to each model θ, satisfying

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

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

We then allow for uncertainty aversion by using a penalization framework based on conditional relative entropy. This framework allows the planner to consider alternative distributions or sets of weights $\tilde{\pi}(\theta)$ across the set of conditional models in a way that is statistically reasonable. This works by restricting the set of alternative model weights considered by the social planner to those that are difficult to distinguish from the prior model distribution using statistical methods. The parameter ξ_a is chosen to determine the magnitude of this penalization. Large values of ξ_a imply low aversion to ambiguity, while small values of ξ_a imply strong aversion to ambiguity. Relative entropy is defined as the expected value of the log-likelihood ratio between two models or the expected value of the log of the Radon-Nikodym derivative between two models.[†]

This new, second-stage problem for the planner is a minimization problem, where the minimization is made over possible distorted probability weights $\tilde{\pi}(\theta)$ which are constrained by ξ_a based on the solutions to the θ conditional value function solutions found previously. This allows the planner to determine the relevant worst-case model for given states of the world to help inform their optimal policy decisions.[‡] 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

$$\tilde{V}_t = \min_{\tilde{\pi}(\theta)} \sum_{\theta\in\Theta}\tilde{\pi}(\theta)(V(\theta) + \xi_a[\log(\tilde{\pi}(\theta)) - \log(\pi(\theta))])$$

$$\text{subject to} \sum_{\theta\in\Theta}\pi(\theta) = \sum_{\theta\in\Theta}\tilde{\pi}(\theta) = 1$$

[†] See (14) 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.

[‡] 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.

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

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

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

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

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

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

4. 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 (15), and the specific algorithm given here builds on the publicly available code and algorithm provided by (7). An online repository with the code, files, and results for our numerical analysis can be found at https://github.com/mbarnet0/Covid.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

$$
\tilde{\beta} = \sum_{\theta} \tilde{\pi}(s_t, i_t; \theta) \beta(\theta), \quad \tilde{\alpha} = \sum_{\theta} \tilde{\pi}(s_t, i_t; \theta) \alpha(\theta),
$$

$$
\widetilde{\beta \zeta} = \sum_{\theta} \tilde{\pi}(s_t, i_t; \theta) \beta(\theta) \zeta(\theta), \quad \widetilde{\beta \zeta^2} = \sum_{\theta} \tilde{\pi}(s_t, i_t; \theta) \beta(\theta) \zeta(\theta)^2
$$

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

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

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

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

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

To solve the HJB equation, we separate the state space into the relevant cases over which we need to impose distinct and necessary conditions for the Markov chain approximation method to solve the HJB equation. We only solve the model for

combinations of the state variables such that $s_t + i_t \leq 1$, thereby ensuring the adding up constraint $s_t + i_t + r_t = 1$ can be satisfied. We begin by first discretizing the values of the continuous-valued state variables in the model as follows:

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

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 $s_s + i_i \leq 1$. Depending on which case we are in, we use a different set of Markov chain Approximation conditions. A visual construction of these conditions in an If-Then code structure is given as follows:

Table S1. If-Then State-Variable Conditions

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

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

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

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

According to the method of (15), the condition we need to satisfy for convergence of our algorithm is that

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

Therefore, we impose

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Because the case with $(\mathcal{R}_0, CFR, \zeta, \alpha) = (1.5, 0.02, 0.65, 0.0)$ produces particularly poor derivative approximations near the boundary, we use the solution to the case with $(\mathcal{R}_0, CFR, \zeta, \alpha) = (1.5, 0.02, 0.35, 0.0)$ as the initial guess for more stable results. For the uncertainty averse cases, we use the uncertainty neutral solution with the same prior weighting as the initial value function guess V_0. This ensures stability in convergence with the additional curvature and decreases the number of iterations required for convergence. Finally, convergence is achieved when the error terms for the policy function and the value function are below given tolerance levels. For the value function error, call it $ValueFunctionError$, we set the tolerance to be $\epsilon_1 = 1\text{e}{-12}$ and use the same criteria as in (7):

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

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

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

where $||\cdot||_2$ represents the 2-norm and again n denotes the iteration number in the Markov chain approximation algorithm. The solutions are qualitatively similar for various different hyper-parameters and initial guesses for which convergence is achieved, and any quantitative differences that arise in a small number of cases are minor. For the most extreme "outside the model" cases with the highest mitigation policy outcomes, i.e., the cases with high infection and fatality rates and low quarantine effectiveness, occasionally the iterative scheme will generate poor derivative approximations near the boundary

which produces unreliable results for some choices of the hyper-parameters. We tested numerous hyper-parameter choices to verify that results are consistent across cases where this does not occur.

Once we have model solutions using our algorithm, results are generated using simulations based on the model solutions where model priors and underlying processes are chosen to fit the scenarios described in the main text. The simulations set shocks to zero, while using solutions assuming there are shocks in the model. Initial values for the simulations are chosen by assuming there is a small initial population of infected individuals, as well as a small initial value of immune/recovered individuals. The first assumption is required to provide a small but rapidly growing level of infection in the model as was observed in the US, and the second is in line with (7) and consistent with some recent research that some people may have at least partial immunity[§], and ensures stability of the simulations that avoid numerical inaccuracies near boundary values as well. The initial values we use are given by $s_0 = 1 - i_0 - r_0$, $i_0 = 0.02$, $r_0 = 0.03$, and $d_0 = 0$.

5. Alternative Prior Weight Scenarios

We provide results for three additional sets of prior beliefs scenarios over the models which we label as follows: (i) "under-estimated," (ii) "over-estimated", and (iii) "split-estimated." Each case highlights important scenarios the planner could confront, and how uncertainty influences policy choices and model outcomes in those different scenarios. For each scenario we show the following results from the uncertainty neutral and the uncertainty averse solutions: Figures S1, S4, and S7 show the planner's optimal quarantine and resulting infections and deaths; Figures S2, S5, and S8 show how the planner's beliefs over $\mathcal{R}_0$, the CFR, the mitigation policy losses, and the mitigation policy effectiveness evolve correspondingly; Figures S3, S6, and S9 show the uncertainty averse planner's distorted probabilities over each model and the detection error probability values for the uncertainty-adjusted solutions.

We first consider the "under-estimated" scenario where the planner assigns a weight of 50% to the model with $\mathcal{R}_0 = 1.5$, initial CFR= 0.005, $\zeta = 0.65$, and $\alpha = 0.0$, with the remaining weight distributed equally to the remaining models. In this case, Figure S1 shows that the uncertainty neutral optimal policy response is quite modest, with quarantine measures starting at zero, peaking around 45%, and finishing after about 12 weeks. This response results in a relatively high number of deaths (almost 2%). The uncertainty averse response significantly increases the initial quarantine level (near 20%), the peak quarantine level (over 60%) and the duration of the quarantine measures (about 20 weeks), drastically reducing the number of deaths. Figures S2 and S3 show that this policymaker substantially decreases weight on the best-case model, transferring much of that weight to the two most-severe models as infections and deaths rise. This leads to implied $\mathcal{R}_0$, CFR, and α that are much higher than originally assumed, and an implied ζ that is much lower than assumed by the prior.[¶] The policymaker reassigns weight to the worst-case outcomes because as the pandemic initially spreads quickly through the population the large number of infected means that underestimating the pandemic becomes even more costly and the planner therefore takes extreme measures to contain the pandemic going forward, approaching policy close to the equal-weighted case in the main text. Thus, in a case where the planner underestimates the severity and economic costs of the pandemic, incorporating uncertainty substantially increases the quarantine policy choice, increasing as the pandemic runs its course, and only shifting back to the prior after a sustained period of low infections.

Next we consider the "over-estimated" scenario where the planner assigns a weight of 50% to the model with $\mathcal{R}_0 = 4.5$, initial CFR= 0.02, $\zeta = 0.35$, and $\alpha = 0.8$, with the remaining weight distributed equally to the remaining models. In this case, Figure S4 shows a higher quarantine policy for the uncertainty neutral solution, with quarantine measures reaching nearly 50% at around 8-10 weeks. The deaths resulting from this response are modest compared to the other scenarios. However, the uncertainty averse response has only a modest impact on the quarantine response, peaking between 55% and 60%, and returning to zero just before 20 weeks. Figures S5 and S6 show that the policymaker maintains high weight on the most-severe model, and has a somewhat higher implied $\mathcal{R}_0$, CFR, and α, and a lower implied ζ. However, given the already high initial weight on the worst-case model, and the fact that the realized outcomes are in fact less severe than what the initial prior weighting would have anticipated, the potential magnitude of probability weights is somewhat limited.

Finally, we consider the "split-estimated" scenario where the planner assigns a weight of 26.67% to the model with $\mathcal{R}_0 = 4.5$, initial CFR= 0.02, $\zeta = 0.65$, and $\alpha = 0.0$, a weight of 26.67% to the model with $\mathcal{R}_0 = 1.5$, initial CFR= 0.005, $\zeta = 0.35$, and $\alpha = 0.8$ with the remaining weight distributed equally to the remaining models. The quarantine, infection, and death outcomes are shown for this case in Figure S7. Given the identical weighted-average parameter values resulting from this case as in the equal-weighted case, the uncertainty neutral results are the same as in that baseline setting. In terms of the uncertainty averse response, we see again that there is a noticeable impact on quarantine measures. However, the effect is more modest in this setting. Figures S8 and S9 show that the policymaker has a fairly high weight on the most severe model, but still maintains a relatively high weight on the least severe model. Thus, the structure of the prior and not just the weighted-value outcomes of the prior, are shown to have some impact. In addition, the implied $\mathcal{R}_0$, CFR, and α, and ζ values show movements similar to the equal-weighted baseline, but are again less significant given this alternative weighting.

The detection error probabilities for each case are shown in Figures S3, S6, and S9. We can see that for each case, we have adjusted the value of ξ_a for each scenario to keep the minimum value roughly consistent across settings. The "under-estimated" case uses an uncertainty parameter value of $\xi_a = 0.0042$, the "over-estimated" cases uses an uncertainty parameter value of $\xi_a = 0.002$, and the "split-estimated" case uses an uncertainty parameter value of $\xi_a = 0.00745$. Each case, apart from the "over-estimated" setting, has a minimum value slightly above 10%. There are a few key differences to note across cases. First,

[§] See articles in Science, Nature, and medRxiv for examples of recent research discussing this.

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

the "under-estimated" and "over-estimated" cases never converge back to the assumed prior weight of 50%, because the model comparisons between uncertainty neutral and uncertainty averse models assume the "under-estimated" and "over-estimated" priors but use the equal-weighted parameter values to generate the model outcomes. As a result, there is always some detection error probability because the assumed model and the data generating model are in fact different. Also, the "split-estimated" converges much more slowly back to the 50% assumed prior weight value because of the more severe divergence in the weighting of the discrete set of model possibilities. Finally, note that the detection error probability does not go below 12%. This is due to the fact that the initial weight on worst-case model is already substantial, limiting potential for probability weight distortions. Combined with the fact that the realized outcomes are in fact less severe than what the initial prior weighting would have anticipated, this limits the magnitude of potential model detection error.

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

Fig. S1. Outcomes With and Without Uncertainty, "Under-Estimated" Prior

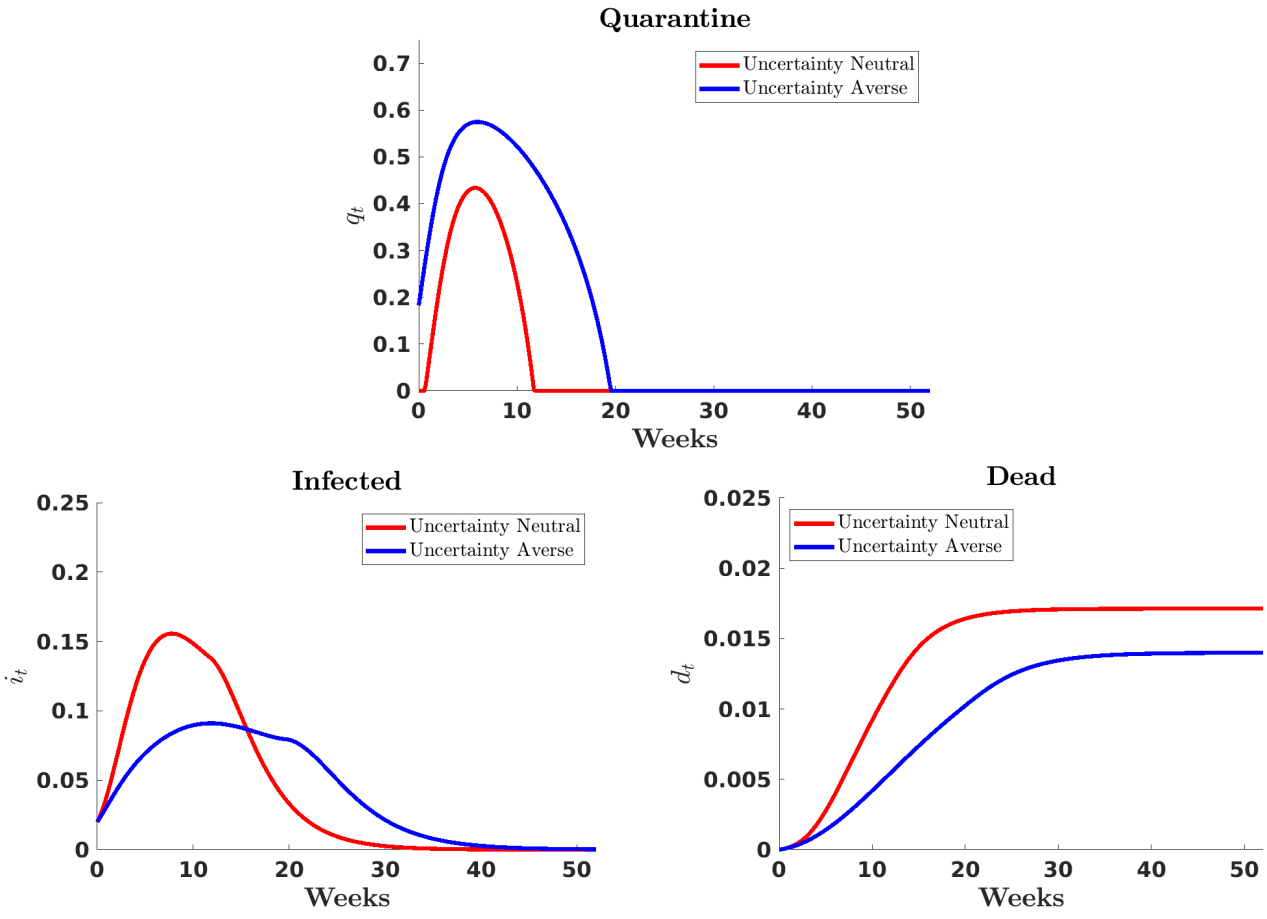

Notes: These figures show the fraction of the population quarantined (top), infected (bottom left), and dead (bottom right) under (i) the uncertainty neutral model response (red) and the uncertainty averse model response (blue).

Fig. S2. Distorted Model Parameters: $\mathcal{R}_0, CFR, \alpha, \zeta$

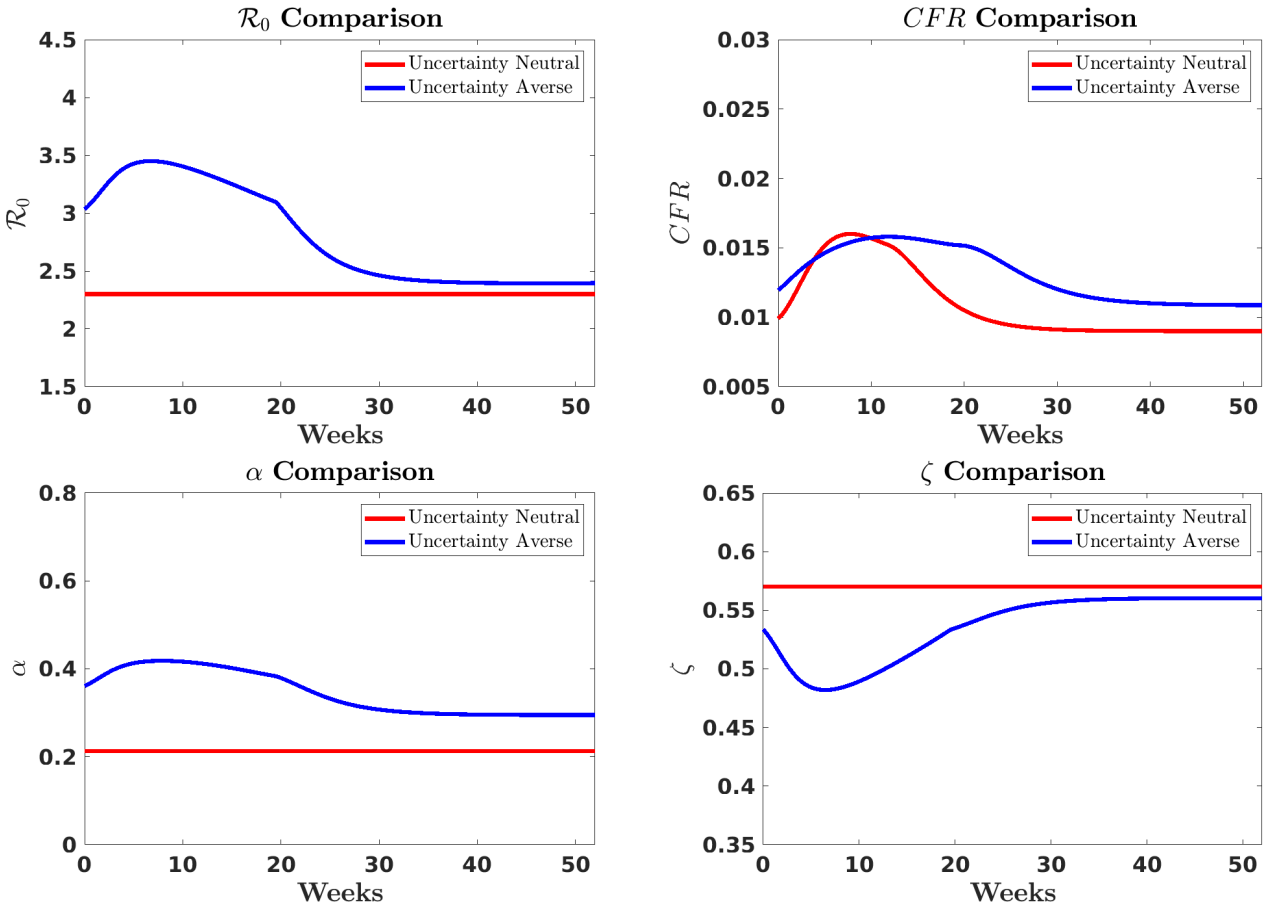

Notes: The figures show the implied $\mathcal{R}_0$ (top left), CFR (top right), α (bottom left), and ζ (bottom right) for the uncertainty neutral case (red lines) and the uncertainty averse case (blue lines). Each case uses the simulation path resulting from its own optimal policy.

Fig. S3. Uncertainty Aversion Distorted Model Probabilities and Detection Error Probability Comparison

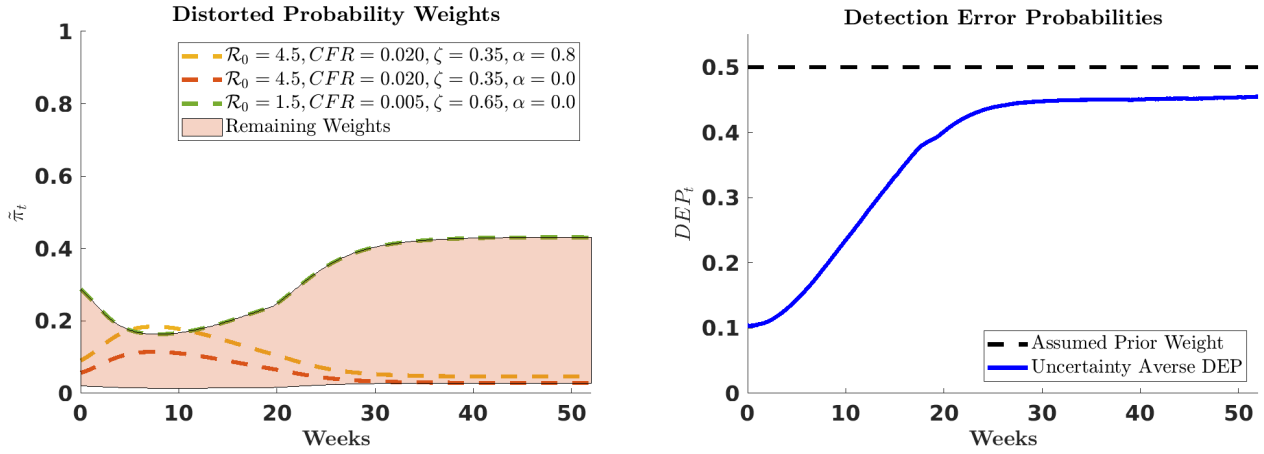

Notes: The left panel shows the distorted probability weights for the uncertainty averse planner. The red dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The yellow dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the highest economic cost of quarantine ($\alpha = 0.8$). The green dashed line is the model with the least severe epidemiological scenario ($CFR = 0.005, \mathcal{R}_0 = 1.5$), the most effective pandemic mitigation scenario ($\zeta = 0.65$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The red shaded area shows the range of distorted probability weights for the remaining models. The right panel shows the detection error probability for model discrimination between the worst-case and baseline models in the uncertainty averse (blue line) setting, as well as the assumed prior weight on the baseline and worst-case models (dashed black line).

Fig. S4. Outcomes With and Without Uncertainty, "Over-Estimated" Prior

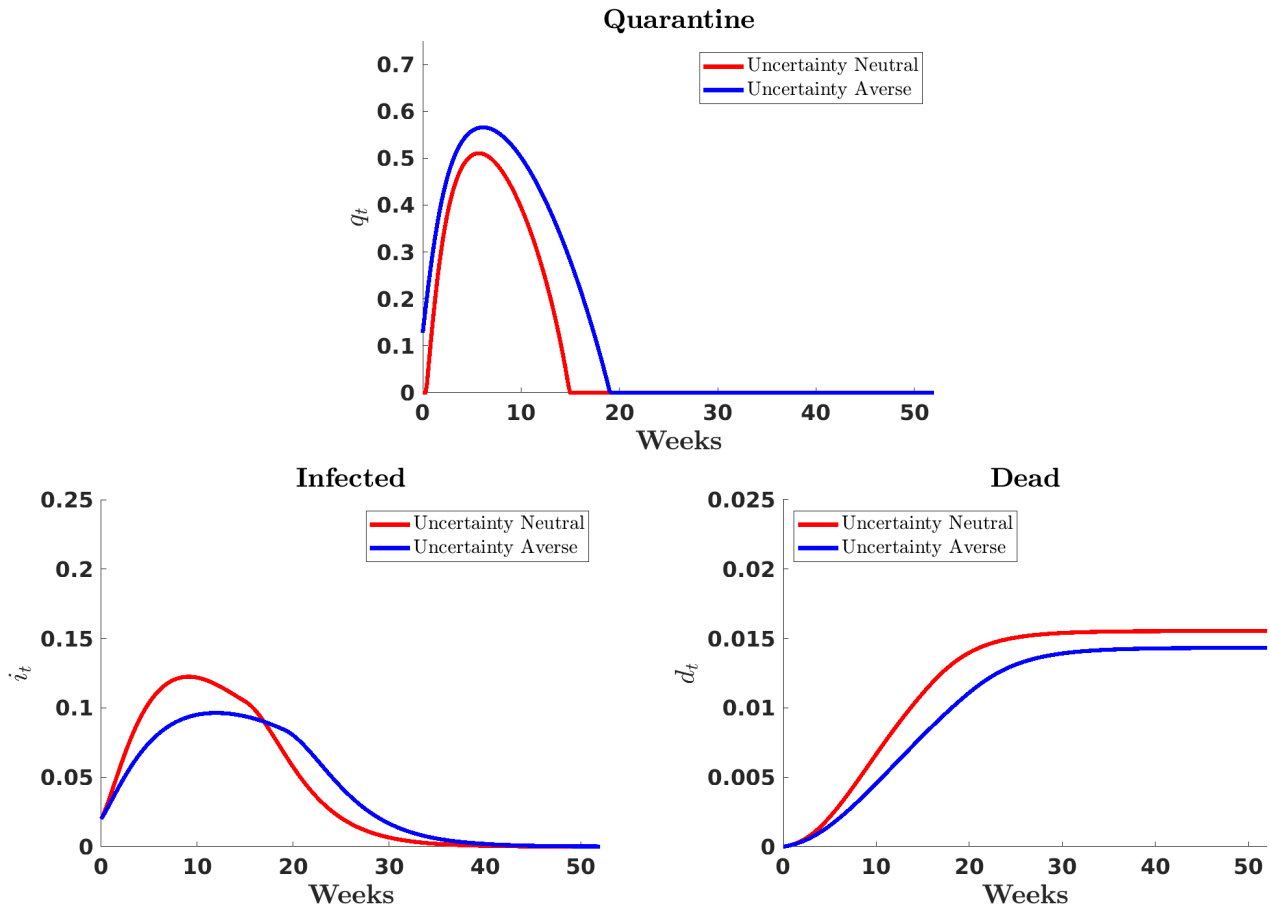

Notes: These figures show the fraction of the population quarantined (top), infected (bottom left), and dead (bottom right) under (i) the uncertainty neutral model response (red) and the uncertainty averse model response (blue).

Fig. S5. Distorted Model Parameters: $\mathcal{R}_0, CFR, \alpha, \zeta$

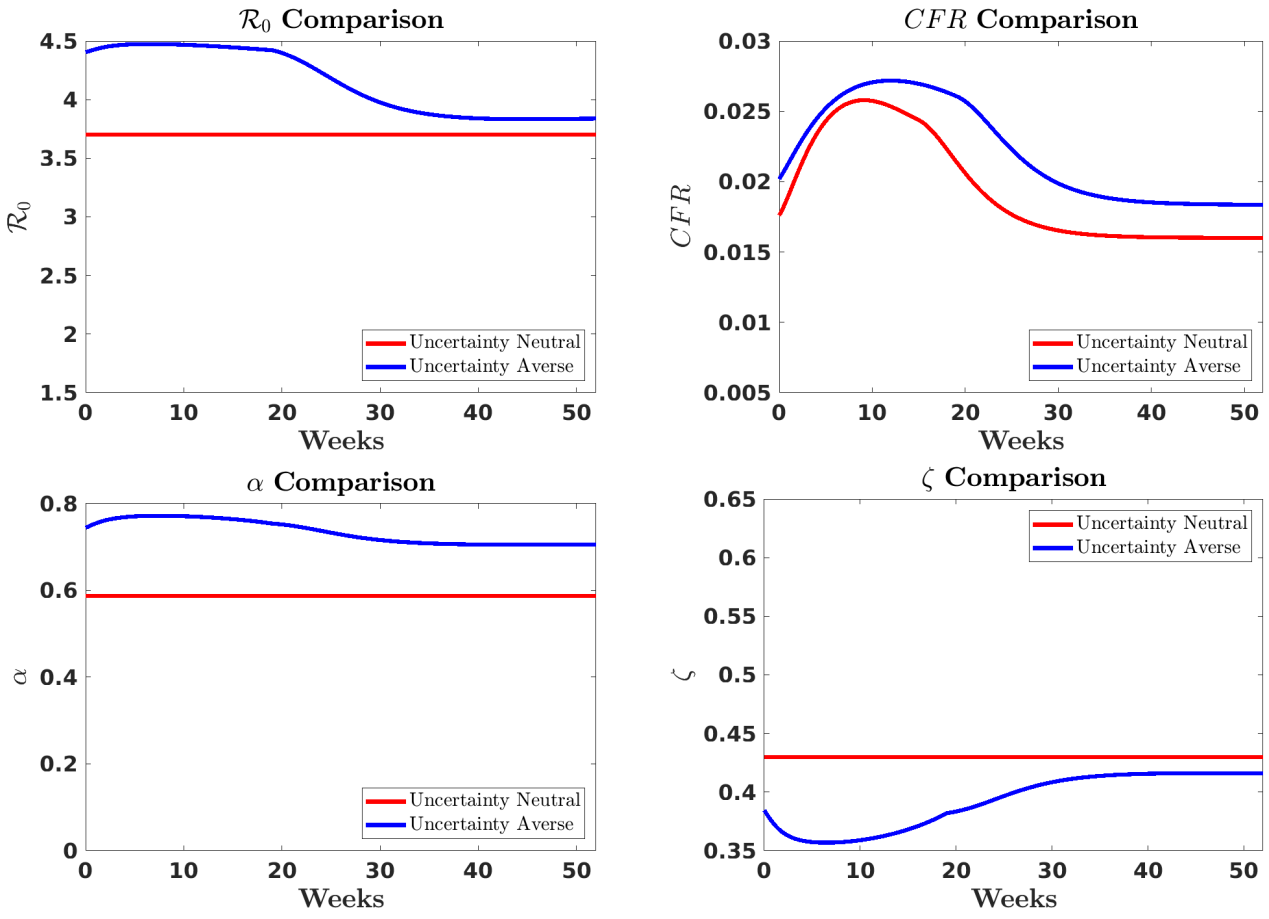

Notes: The figures show the implied $\mathcal{R}_0$ (top left), CFR (top right), α (bottom left), and ζ (bottom right) for the uncertainty neutral case (red lines) and the uncertainty averse case (blue lines). Each case uses the simulation path resulting from its own optimal policy.

Fig. S6. Uncertainty Aversion Distorted Model Probabilities and Detection Error Probability Comparison

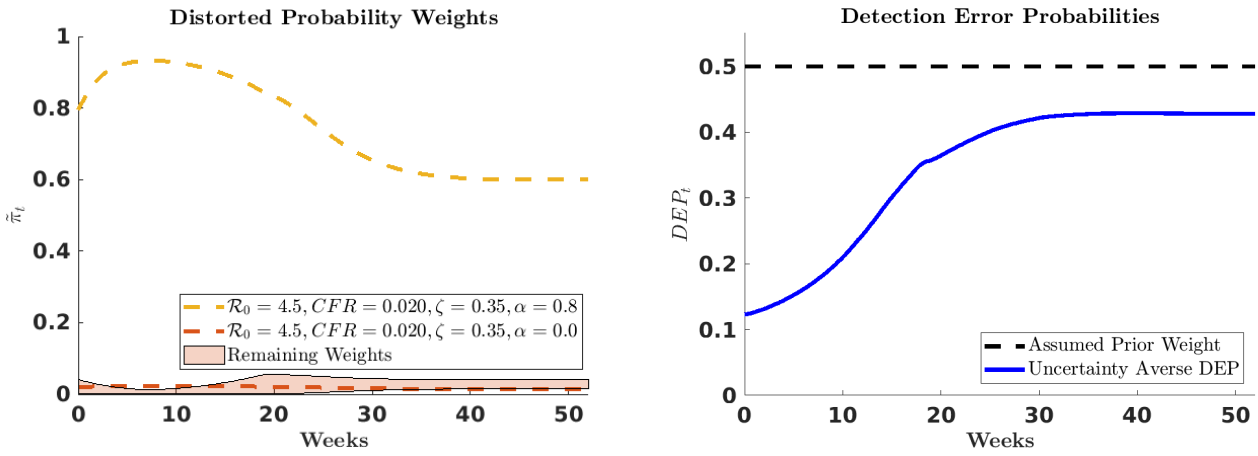

Notes: The left panel shows the distorted probability weights for the uncertainty averse planner. The dashed lines are the probability weights on the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$) with the least effective pandemic mitigation scenario ($\zeta = 0.35$). The red dashed line is the model with the lowest economic cost of quarantine ($\alpha = 0.0$) and the yellow dashed line is the model with the highest economic cost of quarantine ($\alpha = 0.8$). The red shaded area shows the range of distorted probability weights for the remaining models. The right panel shows the detection error probability for model discrimination between the worst-case and baseline models in the uncertainty averse (blue line) setting, as well as the assumed prior weight on the baseline and worst-case models (dashed black line).

Fig. S7. Outcomes With and Without Uncertainty, "Split-Estimated" Prior

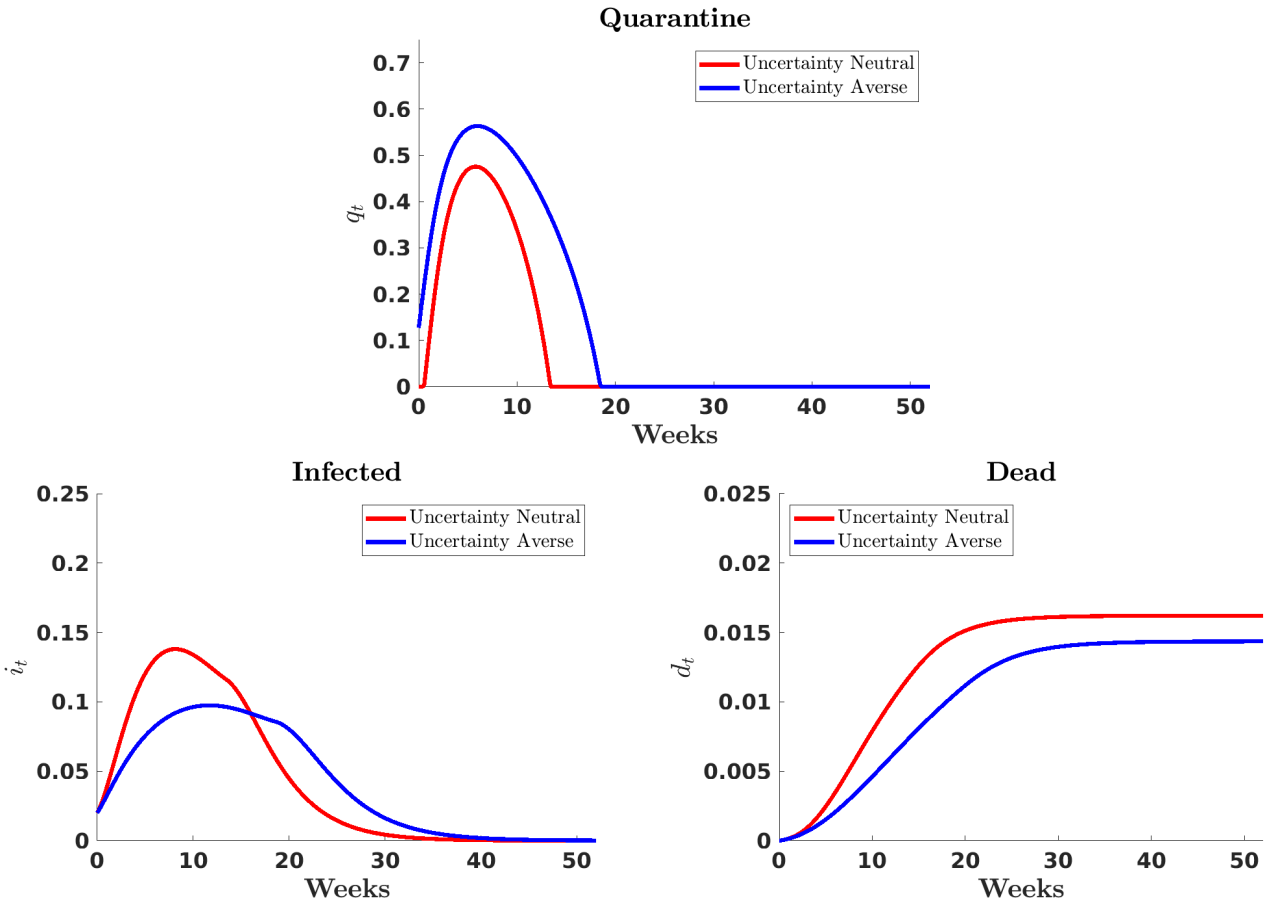

Notes: These figures show the fraction of the population quarantined (top), infected (bottom left), and dead (bottom right) under (i) the uncertainty neutral model response (red) and the uncertainty averse model response (blue).

Fig. S8. Distorted Model Parameters: $\mathcal{R}_0, CFR, \alpha, \zeta$

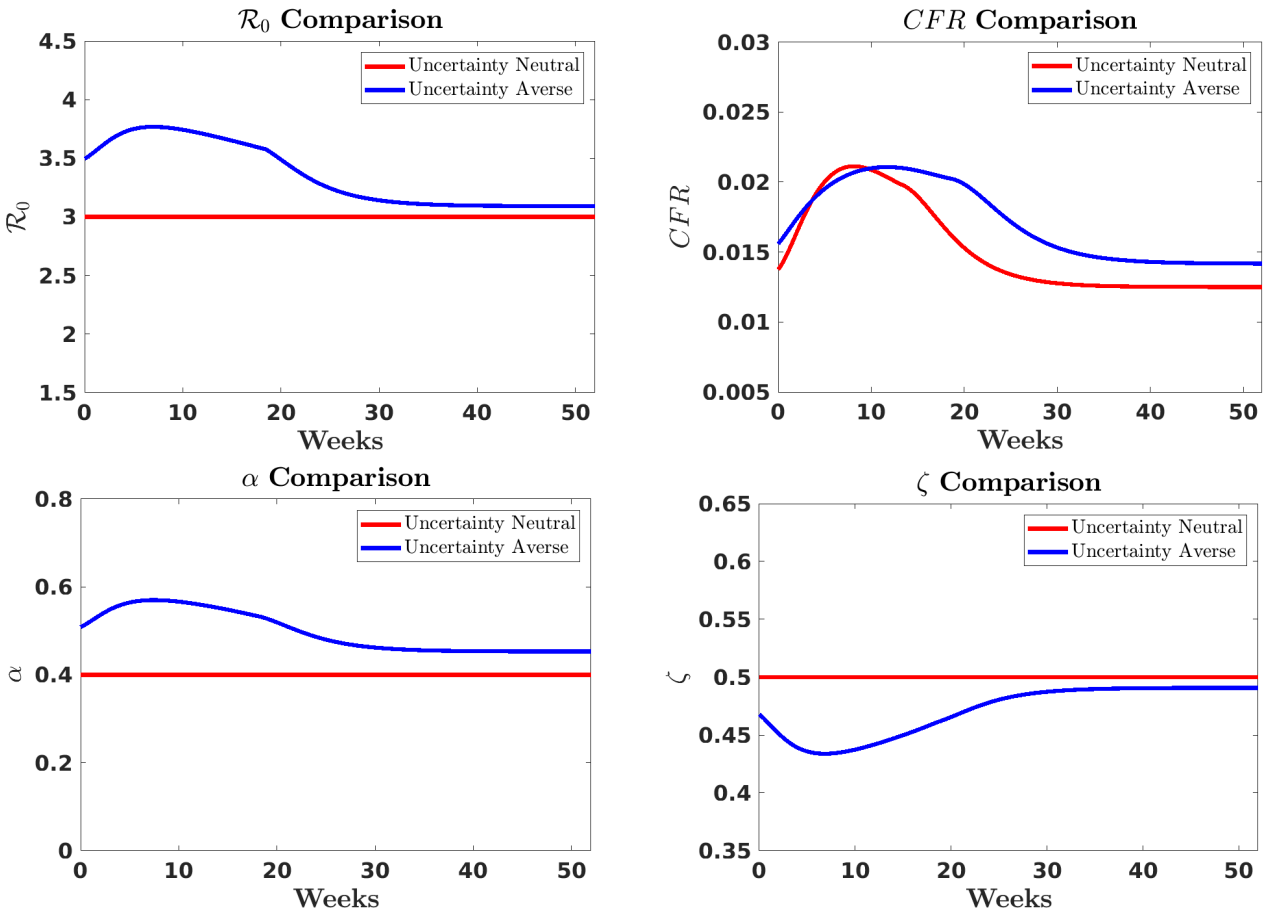

Notes: The figures show the implied $\mathcal{R}_0$ (top left), CFR (top right), α (bottom left), and ζ (bottom right) for the uncertainty neutral case (red lines) and the uncertainty averse case (blue lines). Each case uses the simulation path resulting from its own optimal policy.

Fig. S9. Uncertainty Aversion Distorted Model Probabilities and Detection Error Probability Comparison

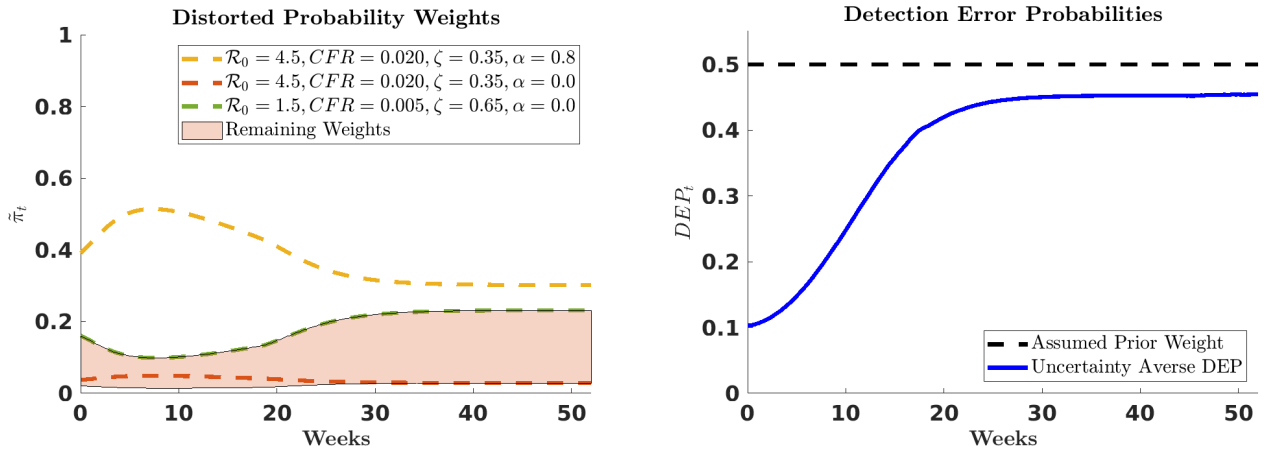

Notes: The left panel shows the distorted probability weights for the uncertainty averse planner. The red dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The yellow dashed line is the model with the most severe epidemiological scenario ($CFR = 0.02, \mathcal{R}_0 = 4.5$), the least effective pandemic mitigation scenario ($\zeta = 0.35$), and the highest economic cost of quarantine ($\alpha = 0.8$). The green dashed line is the model with the least severe epidemiological scenario ($CFR = 0.005, \mathcal{R}_0 = 1.5$), the most effective pandemic mitigation scenario ($\zeta = 0.65$), and the lowest economic cost of quarantine ($\alpha = 0.0$). The red shaded area shows the range of distorted probability weights for the remaining models. The right panel shows the detection error probability for model discrimination between the worst-case and baseline models in the uncertainty averse (blue line) setting, as well as the assumed prior weight on the baseline and worst-case models (dashed black line).

References

1. LP Hansen, TJ Sargent, NE Wang, Robust permanent income and pricing with filtering. *Macroecon. dynamics* **6**, 40–84 (2002).
2. EW Anderson, LP Hansen, TJ Sargent, A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *J. Eur. Econ. Assoc.* **1**, 68–123 (2003).
3. H Chernoff, A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals Math. Stat.*, 493–507 (1952).
4. CM Newman, B Stuck, Chernoff bounds for discriminating between two markov processes. *Stochastics: An Int. J. Probab. Stoch. Process.* **2**, 139–153 (1979).
5. M Barnett, W Brock, LP Hansen, Pricing uncertainty induced by climate change. *The Rev. Financial Stud.* **33**, 1024–1066 (2020).
6. LP Hansen, TJ Sargent, Fragile beliefs and the price of uncertainty. *Quant. Econ.* **1**, 129–162 (2010).
7. FE Alvarez, D Argente, F Lippi, A simple planning problem for covid-19 lockdown, (National Bureau of Economic Research), Technical report (2020).
8. AB Abel, S Panageas, Optimal management of a pandemic in the short run and the long run. *NBER Work. Pap.* (2020).
9. C Jones, T Philippon, V Venkateswaran, Optimal mitigation policies in a pandemic. *Work. Pap.* (2020).
10. M Cagetti, LP Hansen, T Sargent, N Williams, Robustness and pricing with uncertain growth. *The Rev. Financial Stud.* **15**, 363–404 (2002).
11. LP Hansen, TJ Sargent, G Turmuhambetova, N Williams, Robust Control and Model Misspecification. *J. Econ. Theory* **128**, 45–90 (2006).
12. F Maccheroni, M Marinacci, A Rustichini, Ambiguity Aversion, Robustness, and the Variational Representation of Preferences. *Econometrica*, 1447–1498 (2006).
13. Y Izhakian, A theoretical foundation of ambiguity measurement. *J. Econ. Theory* **187**, 105001 (2020).
14. LP Hansen, TJ Sargent, Robustness and Ambiguity in Continuous Time. *J. Econ. Theory* **146**, 1195–1223 (2011).
15. HJ Kushner, P Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*. (Springer Science & Business Media) Vol. 24, (2001).