Overreaction and the Value of Information in a Pandemic *

Keyvan Eslami[1] and Hyunju Lee[2]

[1] *Toronto Metropolitan University, Department of Economics, Email: k.eslami@ryerson.ca*
[2] *Toronto Metropolitan University, Department of Economics, Email: hyunju.lee@ryerson.ca*

November 15, 2022

Abstract

This paper studies optimal mitigation and testing during a pandemic in the presence of partial information. We develop a stylized dynamic epidemiological model where the true number of infected can only be partially inferred from two noisy signals: hospitalization and positivity rate. An egalitarian planner chooses the level of mitigation and testing, respectively affecting the infection rate and signal noise, at some economic cost. We first show that the planner is willing to pay a significant "information premium" to eliminate the uncertainty by extensive testing. However, if testing is prohibitively costly, then a stringent mitigation is optimal, because it partially replaces testing as an information acquisition device. Such policies were often criticized as excessive at the onset of the COVID-19 pandemic. We argue that this "optimal overreaction" is a result of extreme costs of policy mistakes—such as high future casualties—and not an aversion to risk.

JEL Codes: H12, E65, I18

Keywords: Mitigation, Testing, Partial Information, Bayesian Updating, Particle Filtering, COVID-19

*We are grateful to Claustre Bajona, Jesse Matheson, Arthur Sweetman and two anonymous referees for their invaluable comments. We have also greatly benefited from the participants' feedback in the Canadian Health Economics Study Group meetings, 55th Annual Meetings of the CEA, and the Pandemic Workshop at TMU's Department of Economics. The quantitative part of this paper was conducted using the resources of Compute Ontario (computeontario.ca) and Compute Canada (www.computecanada.ca). The authors acknowledge TMU's Faculty of Arts for its financial support. The views expressed are those of the authors.

1 Introduction

Many policymakers, throughout the world, reacted to the advent of COVID-19 pandemic at the starting months of 2020 by pursuing non-pharmaceutical interventions. Such policies ranged from declaring states of emergency, enacting mandatory social distancing measures, locking down parts of economies, to issuing strict stay-at-home orders. Economists, using models that are fine-tuned to match propagation patterns of the disease over the course of the pandemic, have argued that the economic costs of such mitigation policies dominate their benefits in saving lives and livelihoods of people.

Ex post, with enough data on public health policy at different stages of the pandemic and in different economies, there is evidence in support of such claims. However, as we argue in this paper, policymakers' reactions to COVID-19 pandemic cannot be reasonably judged from the lens of historical data that we have after the fact. This is especially true in the earlier stages of the pandemic when every decision was made in an aura of uncertainty. We claim that any evaluation of public health policy in the face of such unprecedented health crises must account for the information limitations that planners may face and, in turn, for the *signalling role* of policy—non-pharmaceutical interventions and testing, in this case.

In order to demonstrate and quantify the potential information value of such policies, we develop a model of disease propagation and mitigation under partial information. We consider a simple neoclassical economy in the context of a compartmental epidemiological model in which otherwise identical individuals are either susceptible to a disease, infected by it, or in the hospitals. We consider the problem of a planner in this setting who seeks to maximize the equally weighted expected lifetime utility of these risk-averse individuals, but cannot directly observe the actual number of infected and susceptible people.

In the model, the planner receives two noisy signals in each period which are correlated with the true number of infected in the economy, namely *hospitalization* and *positivity rate*. While the planner has no control over the standard error of the noise associated with the first signal, she can improve the accuracy of the second signal at a (convex) cost. We also assume that the planner can directly control the infection rate in the economy through an instrument—which we broadly refer to as the *mitigation* policy—at the cost of lowering economic productivity.

Conceptually, one can think of the first signal as the information coming from all sources that the planner has no control over. The most obvious, and perhaps the most important, of such signals is the number of people who go to hospitals in each period (for either out- or in-patient visits as a result of COVID-19 symptoms). As such, while we intentionally try to be stylized in our modelling exercise, we loosely refer to this first signal as the number of hospitalized henceforth.

The second signal consists of any information over which the planner has some control. The most straightforward source of such a signal is testing. Therefore, from now on, we refer to its source as testing and to its realization as the positivity rate.

We quantify our model using data from the province of Ontario, Canada. Using data on hospitalization, positivity rate, and mobility, we estimate a set of parameters, including the variance of signal noise, and eventually the evolution of the infected population based on our model, using *particle filtering*.

In our first set of results, we show that a planner who is only partially informed about the actual number of infected *optimally overreacts* to the pandemic—a term we reserve to indicate a mitigation policy that is more restrictive than the *unconstrained optimum* under full information, due to planner's information constraints. These results hold under reasonable values for initial beliefs and when testing is prohibitively costly—for instance, due to the lack of testing technology or capacity in the earlier stages of the pandemic.[1]

Intuitively, in addition to controlling the spread of the disease, mitigation plays the indirect role of reducing the variance of planner's beliefs *in the future* about the actual number of infected in our environment. That is, mitigation also serves as an *information acquisition device* for the planner. This indirect effect is the main force behind the perceived overreaction in our results. However, the underlying mechanism is not a trivial aversion to risk on the part of the planner. If anything, our quantitative results suggest that the planner's value function is convex in the underlying number of infected, when there is no uncertainty, and, consequently in the priors in the presence of uncertainty.

As such, the planning value, at the beginning of each period, is non-monotone in the variance of priors. The shape of this dependence, however, is highly dependent on the realization of the hospitalized population. We show that a signal that severely contradicts a planner's underlying beliefs entails extreme costs to an uncertain planner. As a result, she tries to avoid such mistakes. Interestingly, when the initial uncertainty is too high but the information content of signals is also high, our planner might decide to delay mitigation until enough signals are received, in an attempt to avoid policy mistakes.

Under our parameter estimates using Ontario's data, our model can account for the policy response to the arrival of first cases in the province reasonably well. Under the assumption that the planner had moderate uncertainty about the actual number of infected, our model suggests that

1. When testing is prohibitively costly, the second signal effectively loses its information content and the economy boils down to one with a single signal—namely, hospitalization.

at least 20% of the initial response could be attributed to the information role of mitigation.

Somewhat expectedly, when the planner underestimates the initial prevalence of the infection, the initial overreaction is up to 15 percentage points more pronounced. What is interesting is that, under our estimated parameters, this initial overreaction is compensated by a rapid reopening of the economy.[2]

To explore the importance of information to the planner further, we consider a counterfactual scenario in which the cost of testing is steep at the onset, but still affordable. Our model suggests that a planner under partial information is willing to give up as much as 17% of output in a single day to eliminate almost all of her uncertainty. This policy is so costly that the planner can no longer afford enacting any mitigation policy at the beginning, leading to a slight increase in cases. However, this is an *information premium* that the planner is willing to pay to avoid policy mistakes.

Relation to the Literature First and foremost, this paper connects the expansive literature on the economics of information and learning to that on the economics of pandemics.

The arrival of the COVID-19 pandemic led to a revival of interest in the economics of pandemics. By combining existing economic models with variations of epidemiological models of disease propagation, many economists have studied the optimal mitigation policies under various assumptions. Two important examples, which this paper builds upon, are Alvarez, Argente, and Lippi (2021) and Jones, Philippon, and Venkateswaran (2021). To the best of our knowledge, however, none of such models consider the importance of uncertainty and learning, on the part of policymakers, at the earlier months of the pandemic.

In this paper, we provide a first-cut attempt at extending this literature to incorporate the role of *ex ante* uncertainty on the optimal choice of policy. In this sense, this paper is best viewed as an investigation of the "value of information," as defined by Gollier (2001), during a pandemic, when policy is made under significant uncertainty.

In this sense, this paper is also related to a literature in epidemiology and economic that has emphasized data inaccuracies and the best choice of signals to guide public health policy (see Brodeur et al. 2021 for an extensive survey), but without models featuring information frictions. By combining the epidemiological model with a rich tradition in macroeconomics, particularly in the study of optimal monetary and fiscal policies under uncertainty (for example, Lucas Jr 1972;

2. These two observations combined can account for Glover et al. (2020)'s criticism of the public health policy in the U.S., which they argue was excessive at the beginning and lifted prematurely.

Townsend 1983; Nimark 2008; Angeletos and Pavan 2009; Melosi 2017), our paper complements this literature by exploring the theoretical implications of such imperfect information on public policy.

In addition, by proposing an estimation method based on our theoretical framework, we contribute to the empirical literature that focuses on using imperfect measures of infection prevalence. Subramanian, He, and Pascual (2021) is just one related example in this literature, estimating the asymptomatic infection using observed cases, based on an epidemiological model.

The literature in the economics of pandemics has emphasized the role of testing as a mitigation measure in itself. *E.g.*, Berger et al. (2022) and Piguillem and Shi (2022) propose epidemiological models with testing and study economic outcomes under different policies. In their papers, tests are used to identify infected individuals in the presence of asymptomatic cases, but there is no learning process about the true number of infected by the planner as in our paper. Chari, Kirpalani, and Phelan (2020), Atkeson et al. (2020), Cleevely et al. (2020), Gollier and Gossner (2020), Eberhardt, Breuckmann, and Eberhardt (2020), Wells et al. (2021) are examples that introduce novel test and quarantine strategies, but their focus is not on the optimal planning problem with partial information as in our paper. Through its emphasis on the significance of partial information, our framework underlines the informational role of testing during an epidemic—a role that has mostly been overlooked before.

The rest of this paper is organized as follows: In Section 2, we present the evidence that has motivated this paper. Section 3 is dedicated to a detailed description of the economy, information structure and the planning problem under consideration. Section 4 presents our quantitative exercise: In section 4.1, we show how we use the model's intuition, together with a novel estimation method, to find the actual number of infected and, in turn, the parameters of the model using data from the province of Ontario, Canada. Sections 4.2, 4.3 and 4.4 solve the model for the estimated parameters under three scenarios: full information; when testing is not available at the beginning; and when testing is available from the onset. We conclude in section 5.

2 A Motivating Observation

After the fact, with enough data on public health policy at different stages of the pandemic and in different economies, there is evidence in support of critics who blame governments for *excessive*

policy reaction to the pandemic at its onset.[3]

In this section, we formally characterize this observation as a motivation for our model and its predictions throughout this paper. Using mobility data from the province of Ontario, Canada, as a proxy for all mitigation measures in place at any point since the onset of the pandemic (February, 2020) to the end of third wave of COVID-19 (June, 2021), we provide evidence in support of excessive application of mitigation policies.[4] We show that the reduction of mobility at the onset of the pandemic was the most severe, even though the hospitalized population and confirmed positives were the lowest compared to the subsequent "waves" of the infectious disease.[5]

The dashed blue line in figure 1 depicts the normalized mobility data, from *Google COVID-19 Community Mobility Reports*, where the pre-pandemic level mobility is normalized to one.[6] The solid orange line in this figure illustrates the hospitalization rate, from the *Government of Ontario*, as a percentage of hospitalized patients with COVID-19 among the populations of ages 30 and over.

As the figure suggests, the mobility measure collapsed sharply to less than 40% of its pre-pandemic level at the very onset of the pandemic in April, 2020, which we know as the *first wave* by now. In the subsequent *second* and *third waves*—marked by an acceleration of infection in December, 2020, and April, 2021, respectively—the mobility measures decreased but not as sharply as during the first wave. While the mobility measure never fully recovered to the pre-pandemic level of one, just before the beginning of the second wave it reached to nearly 80% of the pre-pandemic level. In the subsequent second and third waves, the mobility measure declined by less than 30 and 20 percentage points, respectively, which is in stark contrast to a more than 60 percentage point plunge in mobility during the first wave.

3. In the U.S.,*e.g.*, in a 19-day period between March 19, 2020, to April 7, 42 states issued strict stay-at-home orders (Murray and Murray 2020). In addition, on March 16th, federal guidelines for social distancing for a 15-day period were announced (Barrios and Hochberg 2020). These initial mandates were then gradually lifted—perhaps, as Glover et al. (2020) note, sooner than optimal—never to be enacted again, despite rising case numbers in many states. In the U.K., strict stay-at-home orders were issued on March 23, 2020. These restrictions, however, were almost entirely lifted by July 1, and never enacted again, despite a considerable rise in case numbers as reported by the U.K. government.

4. While not a perfect measure, mobility data has been used extensively in the economic literature as a proxy for all mitigation measures in place in a given economy. For instance, Barrios and Hochberg (2020) and Allcott et al. (2020) use mobility data to document a significant relationship between individuals' partisan views and their adherence to mitigation measures. Simonov et al. (2020), Bursztyn et al. (2020), and Ash et al. (2020) are among examples that document the effect of news outlet exposure on enforceability of mitigation using similar measures.

5. At the time of updating this draft, restrictions imposed to contain the Omicron variant were faced with similar criticisms.

6. We take a simple average of retail and recreation and workplace.

Figure 1. Hospitalization and Mobility in Ontario, Feb, 2020–Jun, 2021

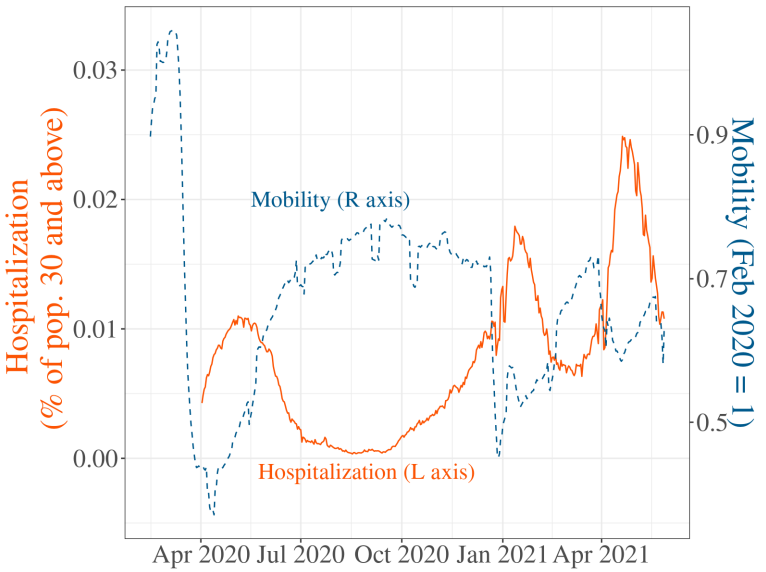

Source: Google COVID-19 Community Mobility Reports and Government of Ontario.

These observations confirm the claims that, compared to the later waves, the province of Ontario pursued and enforced significantly stricter mitigation measures, including a state of emergency,[7] in response to the first wave, when uncertainty was at its heights, even though infection was not nearly as widespread as in the subsequent waves. This excessive reaction—which is in line with, *e.g.*, Glover et al. (2020)'s results—motivates our theory of "overreaction."

In the remainder of the paper, we examine what fraction of this *excessive reaction*—that is a more restrictive mitigation policy than the *unconstrained optimum*, which may or may not have been optimal—can be justified as an *(optimal) overreaction*—which, as mentioned in the introduction, indicates a *constrained optimum*—under differing assumption on the underlying information structure.

3 The Model

Consider an economy in discrete time, consisting of a continuum of individuals of initial measure N_0 at time $t = 0$. People are identical, except for their state of health with regard to a viral disease.

7. As announced by the Government of Ontario on March 17, 2020.

At each date t, of the total population in the economy, N_t, measure S_t are susceptible and at the risk of contracting the disease, and measure I_t are infected.

Absent mitigation, each infected person infects β susceptible individuals on average in each period. On the other hand, infected people recover from the disease at a rate γ per period.[8] Recovery can be in the form of reentering into the susceptible population or succumbing to the disease.[9]

In each period t, fraction θ of the infected population plus a white noise end up in hospitals:

$$H_t = \theta I_t + \varepsilon_t, \tag{1}$$

where $\varepsilon_t \sim \mathbb{N}(0, \nu)$.[10] Of the hospitalized population in period t, h_t, a total of $\varphi(h_t)$ die at the end of period.[11]

Any individual who is not hospitalized produces w units of a perishable consumption good in each period.[12] An individual's period utility from the consumption of this consumption good is

$$u(c) = b + \left(\frac{c^{1-\eta} - 1}{1-\eta}\right), \qquad \eta > 0. \tag{2}$$

An individual's preferences over a stream of consumption take the additively separable expected utility form with discount rate $\rho \in (0, 1)$. We normalize the utility upon death to zero.[13]

8. Note that $1/\gamma$ will be the average duration of disease—in terms of the length of one period—in the model. The ratio β/γ, known as the *reproduction ratio*, characterizes the propagation rate of the disease in society.

9. Unlike different variants of the common epidemiological SIR model that have been widely used in economics, we do not have a "recovered" state in our economy under which individuals are immune to the disease. While there is evidence that individuals recovered from SARS-CoV-2 infection are at the risk of contracting it again, the duration of partial immunity is not yet determined. Though the possibility of instantaneous reinfection is an inaccurate assumption, it serves an important technical purpose in our analysis which will become clearer later on. See appendix D and the end of section 3.1 for a detailed discussion.

10. Even though this formulation is necessary for the analytic tractability of our model, it introduces the possibility of receiving negative hospitalization signals, creating a potential gap between our theoretical and empirical frameworks. While such negative signals never occur in our numerical simulations, as you will see, the mere possibility of such observations can still affect the optimal policy through their effect on planner's beliefs. In our benchmark numerical analysis, we handle this issue by truncating negative hospitalization signals at $H = 0$. As we discuss in our online appendix, a numerical robustness test in which such negative signals are allowed, but the planner never assigns positive weights to negative infection, slightly amplifies our main results.

11. At this point, the only restriction that we impose on $\varphi(\cdot)$ is that $\varphi(h) \leq \gamma h$. If one chooses a convex functional form for $\varphi(\cdot)$, it can capture the capacity constraints on hospitals.

12. We abstract from investment which, given the short period of time under consideration, seems reasonable.

13. The functional form in (2) is the same utility form used by Becker, Philipson, and Soares (2005), Hall and Jones (2007), Ales, Hosseini, and Jones (2014), and Eslami and Karimi (2019) where, intuitively, b captures the "value of being alive," in terms of the utility of consumption per unit time. From a technical standpoint, when η is in $[1, \infty)$,

Finally, we assume the economy starts in period $t = 0$ with a given total and initially infected populations, N_0 and I_0, respectively. We assume that a vaccine is developed in period $T > 0$, when all the infected (and hospitalized) population receive it and the disease is eradicated.

It is worth noting that we can modify this last assumption to capture Alvarez, Argente, and Lippi (2021)'s framework where the cure arrives randomly according to a Poisson process, making the planning problem in the next section stationary and significantly reducing its computational intensity. However, we believe the current setting captures the policymakers' perception of what course the pandemic would take at its onset more accurately, making it a more appropriate framework for the *ex ante* analysis of policy, even if, *ex post*, the COVID-19 pandemic followed a more complicated course.

3.1 Policy and the Information Structure

Our goal is to study the problem of an egalitarian planner in this setting who cannot directly observe the actual number of infected (and susceptible) in $t = 0$ and any subsequent period, but only the number of hospitalized at the beginning of each period. The two policy instruments at the planner's disposal are mitigation and testing.[14,15]

In our environment, mitigation refers broadly to all non-pharmaceutical interventions that can directly affect the *effective* reproduction ratio, by reducing the average number of (transmitting) contacts between individuals. We consolidate all such policies in a variable m. In particular, in each period t, the planner can directly control the infection rate by choosing any m_t in $[0, 1]$. One can interpret $m_t = 1$ as no intervention, whereas $m_t = 0$ indicates the most restrictive lockdown of the economy.

Non-pharmaceutical restrictions—such as social distancing, capacity restrictions or lockdown of various sectors—also reduce the productivity in the economy. This is captured by a non-

$b > 0$ makes sure that the utility of living remains above the utility upon death—for a minimal (sustenance) level of consumption.

14. In what follows, we implicitly assume that the planner can also freely redistribute resources in the economy. This assumption is in contrast to that in, *e.g.*, Glover et al. (2020). For a strictly concave utility function of the form in (2), this implies that the planner equalizes consumption across all living individuals. We will take this trivial result into consideration when writing the planner's objective function.

15. In a robustness exercise, we test the results of our baseline model in section 4.3 when the planner assigns considerable and negative weight to hospitalization. This assumption amplifies the overreaction, when uncertainty is moderate. However, our results regarding under-reaction when uncertainty is high reverses: Now, the planner overreacts even further at the start of the pandemic, instead of waiting for more information.

decreasing function $\Omega : m_t \mapsto \Omega(m_t)$ of the form $\Omega : [0,1] \to [0,1]$, specifying the fraction of economic activity that is allowed under m_t.[16]

Testing in our framework, on the other hand, refers to the integration of all the tools that a policymaker can employ to gain a better insight into the underlying states of the economy. To formalize such function in a stylized way, we assume that, in each period t, and in addition to the hospitalized population H_t, the planner receives an additional signal of the form

$$K_t = \lambda I_t + \upsilon_t(\kappa_t), \tag{3}$$

where λ is the fraction of testing among the infected and κ_t denotes the choice of testing policy in t, taking values in $[0,1]$. In (3), $\upsilon_t(\kappa_t)$ (read *upsilon*) is a white noise of the form $\upsilon_t \sim \mathbb{N}(0, \zeta(\kappa_t))$, where $\zeta(\cdot)$ is assumed to be a strictly decreasing, strictly convex function, with the property that $\zeta(\kappa) \nearrow \infty$, when $\kappa \searrow 0$, and $\zeta(\kappa) \searrow 0$, as $\kappa \nearrow 1$. In this sense, $\kappa_t = 1$ can be construed as testing all the population who are alive for the virus, while $\kappa_t = 0$ means no testing at all—under which, K_t loses all its information content. We will refer to K_t as the positivity rate in t.

The literature on the economics of pandemics has largely focused on testing as an instrument to identify the asymptomatic cases in the population. When coupled with effective quarantining, this is an important channel through which this policy can reduce the *effective* infection rate. We model this direct effect of testing via a strictly decreasing function $\Gamma : [0,1] \to [0,1]$, with the property that $\Gamma(0) = 1$. For any κ, $\Gamma(\kappa)$ characterizes the effect of testing on the infection rate.[17]

Finally, testing is assumed to be costly. In particular, administering κN tests entails a cost of $\Lambda(\kappa N)$ in terms of real output. Function $\Lambda(\cdot)$ is assumed to be strictly increasing and convex, with the property that $\Lambda(0) = 0$.

The timing of the planning problem, given these two instruments, is as follows: At the beginning of period t, the planner receives the news about the realized hospitalization H_t. Next, she decides about the testing policy, κ_t. Given this policy, positivity rate K_t is realized. Finally, she makes a mitigation decision, m_t, based on which production and consumption take place. At the end of the period, $\varphi(H_t)$ of the hospitalized population succumb to the disease.[18]

16. A general functional form for the openness enables us to capture the assumption in Alvarez, Argente, and Lippi (2021) where the planner cannot close the economy below some critical level. In our quantitative exercise, however, we restrict our attention to the case where $\Omega(m) = m$.

17. While stylized, this approach captures the direct effect of testing on the infection rate without introducing an additional state. See Berger et al. (2022) for a model with asymptomatic infection as a state.

18. An alternative timing is for the planner to decide about m_t and κ_t, simultaneously. We briefly discuss the planning problem and the quantitative solution corresponding to this alternative timing in appendix E.

Given m_t and κ_t, the laws of motions of the aggregate states in the economy are as follows:

$$H_t = \theta I_t + \varepsilon_t, \tag{4}$$

$$I_{t+1} = \left[1 + m_t \cdot \Gamma\left(\kappa_t\right) \cdot \beta - \gamma\right] I_t, \tag{5}$$

$$K_t = \lambda I_t + \upsilon_t\left(\kappa_t\right), \tag{6}$$

$$N_{t+1} = N_t - \varphi\left(H_t\right). \tag{7}$$

In addition, given the policy in period t, total population alive, N_t, and total hospitalized population, $H_t = h_t$, the aggregate output in the economy is given by

$$Y_t = \Omega\left(m_t\right) \cdot w \cdot \left(N_t - h_t\right) - \Lambda\left(\kappa_t \cdot N_t\right). \tag{8}$$

In our model, the disease propagates according to (a slight variation of) a SIS model, rather than standard SIR models used by most economists. As you will see in the next section, such an approximation is crucial in keeping our framework analytically and numerically tractable. In appendix D, we demonstrate that, under our parameters estimates, our "linearization" remains an accurate one around the SIR model in the first three months of the COVID-19 pandemic, even in the absence of mitigation. Given that our quantitative results focus on the initial months of the pandemic, we believe that our approximation remains reasonably accurate in addressing the questions that are of interest to this study.

At the end of this section, we must acknowledge that our implicit assumption that the planner knows the characteristics of the disease with complete certainty, while being uncertain about the underlying states, is a rather strong one. The most obvious reason for such an assumption was to keep the complexity of the model at bay. This is a point, however, that we get back to at the end of section 5, and partially address in appendix F.

3.2 The Planner's Problem

Given the information structure, a planner's problem for the economy of section 3 involves solving filtering problems within an otherwise standard dynamic program, in which the planner has to extract information from two separate signals that she receives in each period. In addition, both of these signals are endogenous to the planner's actions.[19]

19. Note that both mitigation and testing policies directly affect the evolution of unknowns in the economy, which, in turn, determine mean signal values. In addition, testing directly determines the noise in the positivity rate.

The timing of the problem, as specified in section 3.1, entails two important implications that greatly simplify the planning problem: First, it allows us to break the planning problem within a period into two "stage" problems—a *first stage* and an *interim stage* problem. Secondly, when it comes to the extraction of signals, even though signals are endogenous to the planner's actions, the timing allows us to *separate* the filtering from optimization problem, *in each stage*.

This "separation principle" allows us to, first, solve the filtering problems to characterize the evolution of planner's beliefs conditioned on the choice of policy and signals' realizations. Afterwards, these beliefs and their laws of motion can be incorporated—as a state and its corresponding control system—into a standard dynamic program for each stage.

Filtering and the Evolution of Beliefs

Our goal is to construct a mean square estimate of some measurable function for each instant t on the basis of a partially observed process. Here, the function under consideration is the planner's value (or continuation value) function, and the partially observed process consists of an underlying process governing the number of infected—the unobserved component—and a signal process—the observed part—both of which are controlled. The optimal filter is characterized by an important result in filtering literature, sometimes referred to as the *representation theorem*.

In discrete time, however, finding the optimal filter boils down to the repeated joint applications of the the *Bayes rule* and laws of motion of states to characterize the evolution of planner's beliefs about the underlying states of the economy—that is the "density" of state variables.

To see how this works, first note that each period t in our economy can be divided into three stages according to the timing within a period: (i) the very beginning of a period, before the hospitalized population is realized (which we denote by superscript *beg*); (ii) after the realization of hospitalization, before a testing decision is made (denoted by *hos*); and (iii) after the positivity rate is realized, before a mitigation decision is made (denoted by *int*).

As elaborated in appendix B, the choice of stages in which we keep track of the planner's beliefs have a critical bearing on the number of states in our problem formulation. Therefore, for the sake of minimizing the set of state variables in the planning problem, we focus on the planner's beliefs in stages (i) and (iii) in each period—right before the realization of H_t and right after the positivity's realization, respectively—and their evolution. If we let $Q_t^s(\boldsymbol{i}) \coloneqq \Pr(I_t \in \boldsymbol{i} \mid \mathscr{I}_t^s)$ be the probability of the infected population being in the set $\boldsymbol{i}$ in stage s, conditional on all the information available, we can denote the beliefs in (i) and (iii) by $Q_t^{beg}(\cdot)$ and $Q_t^{int}(\cdot)$ in period t. We will use $Q_t^{hos}(\cdot)$ to denote the beliefs in stage (ii)—after the realization of hospitalization, before testing decision is made—when needed.

As Kushner and Dupuis (2014) note, this choice of states enables us to break the updating equation into two steps: In the first step, "we update the effects of dynamics," and then we "incorporate the observation" of the signal. The importance of the choice of a linear law of motion for the infected in (5) and linear signal equations of the forms (4) and (6) becomes clear now: If the planner's priors are Gaussian her posteriors will be Gaussian as well, making the evolution of planner's beliefs tractable.

To see how this works, let $Q_t^{beg} \sim \mathbb{N}\left(\mu_t^{beg}, \sigma_t^{beg}\right)$. Then,

$$\Pr\left(H_t \mid I_t = i\right) \sim \mathbb{N}\left(\theta i, \nu\right). \tag{9}$$

As a result, conditioned on observing $H_t = h$ in t,

$$Q_t^{hos} \sim \mathbb{N}\left(\frac{\theta h\left(\sigma_t^{beg}\right)^2 + \mu_t^{beg}\nu^2}{\theta^2\left(\sigma_t^{beg}\right)^2 + \nu^2}, \frac{\nu\sigma_t^{beg}}{\sqrt{\theta^2\left(\sigma_t^{beg}\right)^2 + \nu^2}}\right) =: \mathbb{N}\left(\mu_t^{hos}, \sigma_t^{hos}\right).^{20} \tag{10}$$

Similarly, under the assumption that $\Pr\left(K_t \mid I_t = i, \kappa_t = \kappa\right) \sim \mathbb{N}\left(\lambda i, \zeta\left(\kappa\right)\right)$, and given $Q_t^{hos} \sim \mathbb{N}\left(\mu_t^{hos}, \sigma_t^{hos}\right)$, conditioned on the realization $K_t = k$,

$$Q_t^{int} \sim \mathbb{N}\left(\frac{\lambda k\left(\sigma_t^{hos}\right)^2 + \mu_t^{hos}\zeta\left(\kappa\right)^2}{\lambda^2\left(\sigma_t^{hos}\right)^2 + \zeta\left(\kappa\right)^2}, \frac{\zeta\left(\kappa\right)\sigma_t^{hos}}{\sqrt{\lambda^2\left(\sigma_t^{hos}\right)^2 + \zeta\left(\kappa\right)^2}}\right) =: \mathbb{N}\left(\mu_t^{int}, \sigma_t^{int}\right). \tag{11}$$

Finally, given the law of motion of infected, the choices of policy, $\kappa_t = \kappa$ and $m_t = m$, and $Q_t^{int} \sim \mathbb{N}\left(\mu_t^{int}, \sigma_t^{int}\right)$, we have

$$Q_{t+1}^{beg} \sim \mathbb{N}\left(\left[1 + m \cdot \Gamma\left(\kappa\right) \cdot \beta - \gamma\right]\mu_t^{int}, \left[1 + m \cdot \Gamma\left(\kappa\right) \cdot \beta - \gamma\right]\sigma_t^{int}\right) = \mathbb{N}\left(\mu_{t+1}^{beg}, \sigma_{t+1}^{beg}\right). \tag{12}$$

Intuitively, the updating equations in (10) and (11) characterize the optimal filter by specifying the "optimal weight" the planner must assign to her priors and to the signal. This is most clearly observed if we write μ_t^{hos} and μ_t^{int} as

$$\mu_t^{hos} = \frac{\theta h\left(\sigma_t^{beg}\right)^2/\nu^2 + \mu_t^{beg}}{\theta^2\left(\sigma_t^{beg}\right)^2/\nu^2 + 1} \qquad \textit{and} \qquad \mu_t^{int} = \frac{\lambda k\left(\sigma_t^{hos}\right)^2/\zeta\left(\kappa\right)^2 + \mu_t^{hos}}{\lambda^2\left(\sigma_t^{hos}\right)^2/\zeta\left(\kappa\right)^2 + 1}.$$

20. This is derived by substituting (9) into the updating equation—equation (26) in appendix B.

These equations show that, the lower the variance of the noise, the more emphasis the planner has to put on her new observations, and less on her past priors.

A Two-Stage Planning Problem

We can now write the planner's problems in the first and interim stages, replacing the filtering problem by the updating equations, equations (10), (11) and (12).

To this end, let us first combine (10) and (11) to write planner's posteriors in some period t, after $K_t = k$ and $H_t = h$ are realized—*i.e.* $Q_t^{int}(\cdot)$—as a function of her priors at the very beginning of the period—that is $Q_t^{beg}(\cdot)$—given the choice of testing κ:

$$\mu_t^{int}(K_t = k) = \left[\frac{\lambda k s_1^2 + s_2 \zeta(\kappa)^2}{\lambda^2 s_1^2 + \zeta(\kappa)^2}\right] \quad \textit{and} \quad \sigma_t^{int} = \left[\frac{\zeta(\kappa) s_1}{\sqrt{\lambda^2 s_1^2 + \zeta(\kappa)^2}}\right]. \tag{13}$$

where $s_1 := \nu \sigma_t^{beg} / \sqrt{\theta^2 \left(\sigma_t^{beg}\right)^2 + \nu^2}$ and $s_2 := \left[\theta h \left(\sigma_t^{beg}\right)^2 + \mu_t^{beg} \nu^2\right] / \left[\theta^2 \left(\sigma_t^{beg}\right)^2 + \nu^2\right]$. In equation (13), μ^{int} is written as a function of K to emphasize the fact that, *ex ante*, it is a random variable whose realization depends on the realization of the positivity rate. From the perspective of the planner's beliefs, the distribution of this random variable is as follows:

$$\mu_t^{int} \sim \mathbb{N}\left(\mu_t^{hos}, \frac{\lambda \left(\sigma_t^{hos}\right)^2}{\sqrt{\lambda^2 \left(\sigma_t^{hos}\right)^2 + \zeta(\kappa)^2}}\right), \tag{14}$$

where μ_t^{hos} and σ_t^{hos} are given in (10).

Then, suppressing time-subscripts when possible for the sake of notational brevity, the planner's first-stage problem in period $0 \leq t < T$, given the population alive (N), priors before updating based on the hospitalized realization (Q^{beg}) and the realization of hospitalization ($H = h$), can be written recursively as

$$\begin{aligned} V_t\left(N, H = h, Q^{beg} \sim \mathbb{N}\left(\mu^{beg}, \sigma^{beg}\right)\right) \\ = \max_{\kappa \in [0,1]} \mathbb{E}\left[V_t^{int}\left(N, H = h, Q^{int} \sim \mathbb{N}\left(\mu^{int}(K), \sigma^{int}\right), \kappa\right)\right], \end{aligned} \tag{15}$$

subject to laws of motion of beliefs in (13). In this problem, $V_t^{int}(\cdot, \cdot, \cdot, \cdot)$ is the planner's interim value in period t. The expectation in problem (15) is evaluated with respect to μ^{int}, given its distribution in (14).

The planner's recursive interim-stage problem, for any $0 \leq t < T$, is:

$$V_t^{int}\left(N, H = h, Q^{int} \sim \left(\mu^{int}, \sigma^{int}\right), \kappa\right) = \tag{16}$$
$$\max_{m \in [0,1]} \left\{ N \cdot u\left(Y/N\right) + \rho \mathbb{E}\left[V_{t+1}\left(N - \varphi\left(h\right), H', \left(Q^{beg}\right)'\right)\right] \right\}$$
$$s.t. \quad Y = \Omega\left(m\right) \cdot w \cdot \left(N - h\right) - \Lambda\left(\kappa \cdot N\right), ^{21}$$

where $\left(Q^{beg}\right)'$ is specified by (12).

The expectation in (16) is an iterative expectation operator, first with respect to the conditional distribution of H', given I' (as in (9)), and then with respect to the distribution of I', as perceived by the planner. However, we can use the law of iterated expectations to write this as the expectation with respect to the unconditional distribution of H'. Given planner's priors at the beginning of period $t+1$ about the distribution of infected, $\left(Q^{beg}\right)'$, this unconditional distribution is

$$H' \sim \mathbb{N}\left(\theta\left(\mu^{beg}\right)', \sqrt{\theta^2\left[\left(\widetilde{\sigma}^{beg}\right)'\right]^2 + \nu^2}\right). \tag{17}$$

In the terminal date ($t = T$) and after the introduction of the vaccine, planner's value becomes

$$V_T\left(N, H = h, Q^{beg} \sim \mathbb{N}\left(\mu^{beg}, \sigma^{beg}\right)\right) = N\left[u\left(\left(\frac{N-h}{N}\right)w\right) + \rho\frac{u\left(w\right)}{\left(1-\rho\right)}\right]. \tag{18}$$

The Unravelling of Uncertainty

Before discussing our solution method and demonstrating our quantitative results, it is worthwhile to examine the implications of the updating equations for the optimal policy. A critical observation with regard to mitigation in our economy is its impact on uncertainty. As equation (12) shows, the variance of posteriors decreases as the planner imposes more restrictions on the economy. This observation suggests that, other than restricting the spread of the disease, mitigation also serves as an important *information acquisition device*.

Next, note that, as $\kappa \to 0$ in (11), Q_t^{int} converges to Q_t^{hos} (in a strong sense). Intuitively, this means the positivity rate loses its information content as fewer people are tested in a period. It is informative to compare the case of $\kappa = 0$ with one where $\kappa \to 1$. In this latter scenario, the variance of μ^{int} in (14) tends to σ^{hos}. It is important to remember that this is different from the planner's posterior *after* observing K. When $\kappa = 1$, the planner will have *full information* about the state of the economy *after* observing K. (This is clear from equation (11), with $\sigma^{int} \to 0$ as $\kappa \to 1$.) However, her prior about "which" value of K will be realized at the time of choosing

$\kappa = 1$ is by no means certain. This prior is determined by her beliefs about the distribution of infected in the current period—*i.e.* by $\mathbb{N}\left(\mu^{hos}, \sigma^{hos}\right)$.

The above argument underlines two tensions in the planning problem when choosing the testing policy: If, we assume $\Gamma(\kappa) = 1$ for the sake of argument, an increase in κ, on the one hand, leads to a decline in the next period's variance of beliefs, and, on the other hand, makes the information content of K more valuable. This means that the planner's uncertainty increases in the immediate future. In other words, an increase in κ "brings uncertainty from the future to the present." As we will discuss in section 4, a planner's willingness to do this is driven by two forces: An inclination for ignorance, and an aversion to making policy mistakes. We will return to this point later on.

4 Quantitative Results

The functional equations in (15) and (16) characterize a non-stationary dynamic program which is analytically intractable. However, one can use a standard iterative method to solve this program numerically, to examine the properties of its solution and the policy functions associated with it. This is done by starting from the terminal condition (18), and solving for $V_t(\cdot,\cdot,\cdot)$ and $V_t^{int}(\cdot,\cdot,\cdot,\cdot)$ as functions of $V_t^{int}(\cdot,\cdot,\cdot,\cdot)$ and $V_{t+1}(\cdot,\cdot,\cdot)$, respectively, for a given choice of functional forms and parameter values.[22]

To this end, we choose one day as the length of one period in our model, and we assume that the vaccine is developed after 540 periods (18 months). As mentioned at the beginning of section 3, though far from a consensus, this was the most widely-held belief among the public about the course of the pandemic at its onset. As such, we find it a reasonable benchmark to capture the planner's beliefs at the start of the pandemic.

We normalize N_0 to 100, and set $\eta = 2$ in (2) as is customary in the macroeconomic literature. The annual discount rate is set to 0.95 (to incorporate the natural mortality rate). Value of being alive, b in (2), is chosen such that the value of statistical life is equal to \$10m.[23]

The functional form of $\zeta(\kappa)$, the standard deviation of testing, is set so that it is strictly decreasing and convex with the testing policy κ. The standard deviation parameter ν_κ is normalized with

22. Even from a numerical standpoint, the problems in (15) and (16) are extremely resource-intensive and time consuming: These are non-stationary problems in a minimum of five state variables. In addition, there is no guarantee that they are concave programs. We discuss the novel solution method used for the simulations in the online appendix.

23. This is the midpoint of statistical value of life estimates in the literature, and at the upper bound of the value suggested by Hall and Jones (2007).

the initial testing rate $\kappa_{init} = 4.08 \cdot 10^{-4}$.

$$\zeta(\kappa) = \kappa_{init} \cdot \nu_\kappa \cdot \left(\frac{1}{\kappa} - 1\right). \tag{19}$$

The fatality function is

$$\varphi(H) = \min\left\{\varphi_1 H + \varphi_2 H^2, H\right\}, \tag{20}$$

where φ_1 and φ_2 are calibrated using Government of Canada's cumulative hospitalization and fatality reports by July 9, 2021.[24]

We assume an average Canadian earns $125 a day, and choose the following functional form for the cost of testing:

$$\Lambda(\kappa \cdot N) = a_1 \cdot \kappa \cdot N + a_2 \cdot (\kappa \cdot N)^2. \tag{21}$$

In section 4.3, a_1 and a_2 are chosen prohibitively large, to capture the more relevant scenario in which testing capacities are fully developed only later in the pandemic. In sections 4.2 and 4.4, under the counterfactual that testing is available from the beginning of the pandemic, a_1 is chosen such that a single test entails a flat cost of $50, and a_2 such that testing 10% of population in a single day calls for 10% of the daily output.

Finally, we let $\Omega(m) = m$ and

$$\Gamma(\kappa) = \frac{1}{(1 + e \cdot \kappa)}. \tag{22}$$

We choose e such that, at its full capacity ($\kappa = 1$), testing can reduce the infection rate by 50%.

4.1 Estimation and Data

We use two main datasets to estimate the epidemiological parameters of the model—*i.e.* β, γ, θ, λ, ν, ν_κ and I_0: COVID-19 epidemiological data including hospitalization, deaths, confirmed positives, and number of tests administered, and mobility data in Ontario, Canada, from April 2, 2020, to May 28, 2021.[25] Our epidemiological data includes the number of patients hospitalized and deceased with COVID-19, number of confirmed positive tests, and the total number of tests sourced from the Government of Ontario. Mobility data comes from Google COVID-19 Commu-

24. This is similar to the fatality function used by Alvarez, Argente, and Lippi (2021).

25. April 2, 2020 is the earliest date on which the hospitalization data is available for the province of Ontario.

Figure 2. Hospitalization and Mobility in Ontario

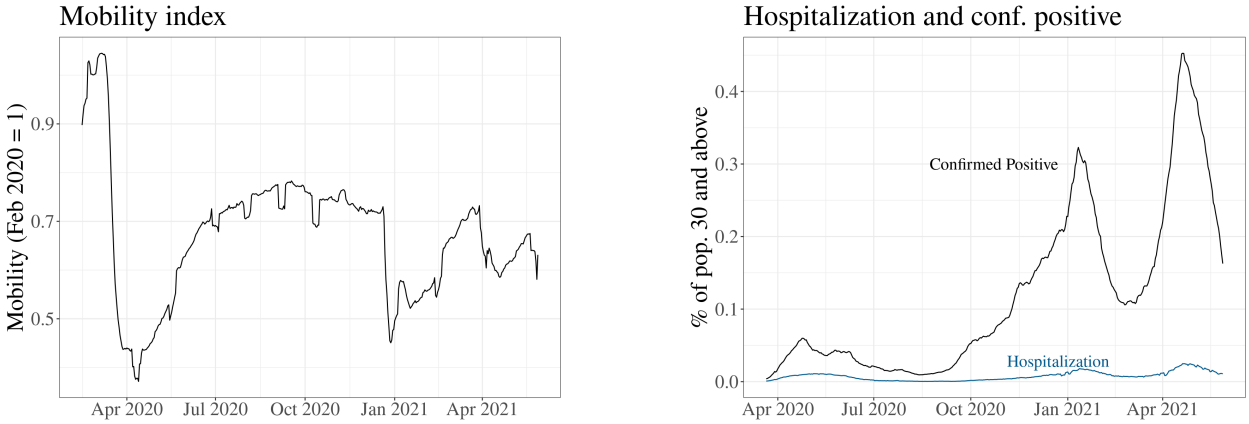

Source: Google COVID-19 Community Mobility Reports and Government of Ontario.

nity Mobility Reports, for the province of Ontario, Canada.[26]

Figure 2 illustrates the hospitalization rate (solid blue line) and confirmed positives (solid black line) as a fraction of the population 30 years old and over on the right panel, and the mobility index on the left panel, where the average mobility in February, 2020, is normalized to 1. More details of the data construction and sources can be found in appendix A. Model counterparts for the hospitalization, confirmed positives, and mobility are H, K, and m, respectively, as given by equations (4)–(6). Using the changes in mobility index compared to the pre-pandemic period as a proxy for the degree of mitigation m, which is the control variable in the interim planning problem in (16), is a widely used exercise in the literature (e.g., see Bargain and Aminjonov 2020; Barrios et al. 2021). Finally, the number of deaths due to COVID-19 accounts for the changes in the total population (N) in the model.

In using the data for calibration, we assume that the hospitalization and confirmed positive data series are *lagged* signals of the true number of infected I. More specifically, we assume that hospitalization and confirmed positive data is a signal for the true number of infected 35 days prior to the observation. This is motivated by a seemingly negative correlation between the mobility rate and hospitalization and confirmed positive rate, as can be observed in figure 2. For example, by the end of December, 2020, the mobility index dropped sharply as the "strict" miti-

26. As Brodeur et al. (2021) argue, there are many advantages but also shortcomings in using the mobility data. One important note is that mobility data combines all the mitigation policies in effect, from social distancing to lockdown and self-isolation. The potential downside is that it also involves behavioural factors. In the context of our model, our assumption is that a planner is already internalizing these behavioral factors and suggesting a decentralized policy that brings about the unconstrained optimum. However, we acknowledge that the mobility data does not provide situation context for the data reported.

Parameter	Value	Description	Source
γ	1/18	Recovery rate	Alvarez, Argente, and Lippi (2021)
θ	0.025	Hospitalization rate	CDC
λ	0.302	Fraction of testing	θ/ frac.hospitalized*
ν	0.036	Signal noise std, hospitalization	Joint estimation
ν_κ	0.814	Signal noise std, testing	
β	0.091	Infection rate	
δ	0.176	Variant infection factor	
I_0	0.071	Initial infected, percentage	

Table 1. Calibration of Epidemiological Parameters

gation measures were implemented in the province. However, the hospitalization and positivity rate increase sharply *after* the provincial shutdown has begun. Under our benchmark structure, where the strict mitigation measures (low m) suppresses the number of infected (I) immediately without delay and both hospitalization (H) and confirmed positive (K) are meaningful signals for the infected, such observation of strict mitigation measures and simultaneous spike in hospitalization and confirmed positive is extremely unlikely through the lens of the model. In order to reconcile these seemingly counterintuitive observations and maintain the tractability of the model, we assume that hospitalization and confirmed positive data at time t are signals for the number of infected in $(t-35)$, where one period corresponds to one day.[27]

We use these two data series of hospitalization/confirmed positivity rate and mobility index to estimate the standard deviation of hospitalization signal noise (ν), testing signal noise (ν_κ), the infection rate (β), and the initial percentage of infected (I_0), using *particle filters*.[28] In order to account for the Alpha variant, which was first detected in Ontario in January, 2021, and is known to be more contagious than the previous strains (Jüni et al. 2021), we estimate a "variant infection factor" parameter, δ, which augments the infection rate from January 1, 2021, to the end of the sample period.[29] Our estimation method involves maximizing the likelihood of our infection control system, using a sequential Monte Carlo (particle filter) algorithm designed for non-linear partially-observed Markov processes (POMP) (King, Nguyen, and Ionides 2016). The details of

27. Admittedly, the exact number of days in delay (35 in this case) is difficult to measure. We have matched the peak of the second wave hospitalization date and the sharp drop in mobility in this calibration. This is designed to give a conservative measure of the key parameters ν and ν_κ, the standard deviations of hospitalization and confirmed positive signal noise, respectively.

28. Since the N_0 is normalized to 100, I_0 is the percentage of population 30 and over that are infected at time $t = 0$.

29. In other words, β in equation (4) has now an additional coefficient: $\beta(1 + \delta \mathbb{1}_{\{year=2021\}})$.

Figure 3. Hospitalization Data vs Model Simulation Results

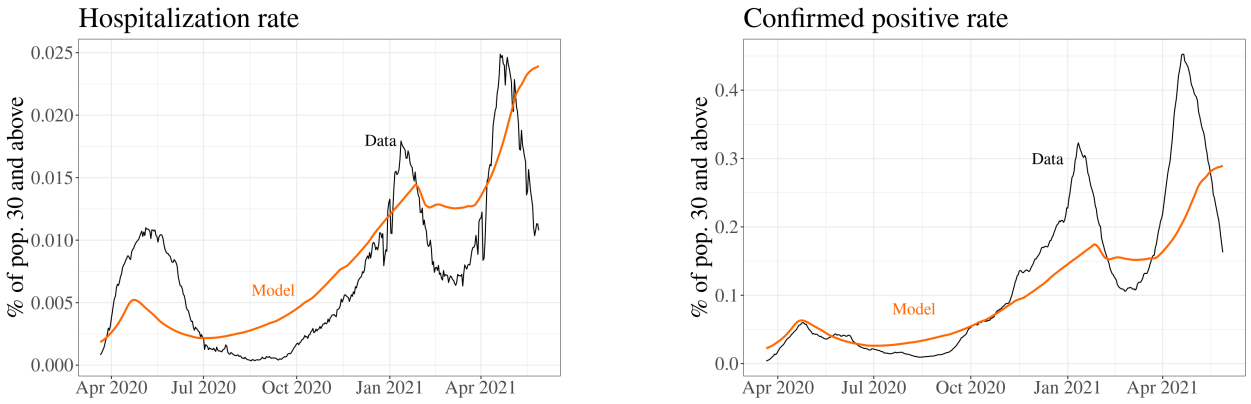

Source: Government of Ontario and authors' calculations.

particle filtering algorithms, including the settings for the method of iterated filtering and *pomp* R-package used in the calculation, are discussed in appendix C.

Two of the three remaining epidemiological parameters, namely the recovery rate (γ) and hospitalization rate (θ), are borrowed from the literature and the estimations of CDC, respectively. Finally, the fraction of testing among the infected (λ) is calculated as the ratio of hospitalization (θ) over the fraction of hospitalized among the confirmed positives in the data, which is on average equal to 0.083 over the sample period.[30] Table 1 summarizes the estimation results.

Using the estimated parameters, we can compare the hospitalization and confirmed positives data and the model predictions in order to evaluate the fitting of our epidemiological model to the data. Figure 3 depicts the data hospitalization rate (solid black line in the left panel) overlaid with the simulated mean hospitalization rate (solid orange line in the left panel). Analogously, the positivity rate (solid black line in the right panel) and its simulated mean model counterpart (solid orange line in the right panel) are compared. The model performs particularly well in predicting the second wave (with its peak hospitalization rate around mid-January, 2021), while the first wave peak hospitalization rate in the model is earlier than that of the data. The predicted hospitalization peak during the third wave lags that in the data. Overall, given the restrictions of the model parameters—for example, a fixed recovery rate (γ), which, in reality might have changed over the course of pandemic depending on variants or hospital capacities—our simplified framework can fit into the recurring waves of COVID-19 outbreaks in Ontario rather well.

In the remainder of this section, we use our results in table 1 to solve and compare the planning

30. This follows from the fact that the average number of hospitalized we observe in the data is also a fraction of confirmed positive cases.

problem, numerically, under different scenarios.

4.2 Case of Full Information

Following the standard practice in the information economics literature, we begin our analysis by examining the case of full information—that is, when the planner knows the underlying state of the economy with certainty. This also allows us to test our "overreaction" hypothesis later on, by comparing the optimal policy with and without uncertainty—*i.e.*, constrained vs unconstrained optima.

The assumption of full information corresponds to a case where $\sigma^{beg} = 0$ in (15) and (16). When that is the case, K loses its information value, and we can combine these two problems as

$$V_t(N, H = h, I) = \max_{m,\kappa \in [0,1]} \left\{ N \cdot u(\Omega(m) w(N-h)/N) + \rho \mathbb{E}\Big[V_{t+1}(N - \varphi(h), H', [1 + m \cdot \Gamma(\kappa) \cdot \beta - \gamma] I) \Big] \right\}, \quad (23)$$

subject to the same terminal condition as before. In equation (23), we have replaced Q^{beg} by I, as the planner's belief about the average number of infected people coincides with the actual infected population. The expectation in (23) is with respect to the (conditional) distribution of H' in (9).

Figure 4 depicts the solution to the Bellman functional equation in (23) for $t = 0$ and two different values of H_0, as a function of I_0, when testing is available from the onset at the cost given in (21). While far from reality, this counterfactual helps highlight the mitigation role of testing in our setting when information is not a concern.

As one can expect, the average lifetime expected utility is decreasing in the initially infected population.[31] An uncommon feature of this value function, however, is its convexity in the number of infected. This implies, if everything else could be kept unchanged, a mean preserving spread of the beliefs around some value of I_0 would in fact increase the planner's expected payoff—that is the planner would prefer uncertainty over the average number of infected.

Figure 5 shows the sample path of optimal mitigation policy in the first 310 periods, and the re-

31. While the numerical value of the function is not relevant, it is worth mentioning that the value function itself remains positive for all values of I_0 of interest—that is above the value upon death. If this were not the case, the disease would create a possibility for the planner to kill as many people as possible!

Figure 4. Initial Value as Function of Infected Population

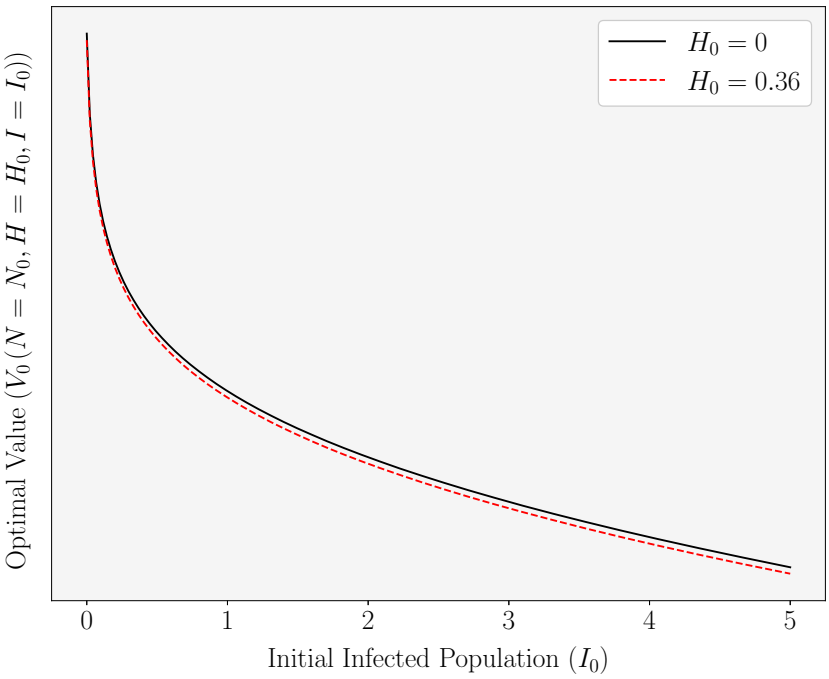

sulting paths of infected, (accumulated) deceased, and hospitalized population under the optimal policy. Each sample path in each panel corresponds to one level of initial infection.

In all cases, optimal policy involves no testing in the first 310 periods ($\kappa_t = 0$). Therefore, according to our model, controlling the spread of the disease must be entirely delegated to non-pharmaceutical mitigation, when information is not a concern.[32]

As panel (d) of figure 5 shows, the optimal initial reaction to the pandemic becomes more strict as the number of initial cases increases, as expected. However, these initial differences in optimal mitigation are neither significant nor prolonged. What stands out is the time it takes for the restrictions to be lifted as the number of initial cases increases, ranging from 50 days to about 200 in figure 5. In other words, optimal mitigation under full information mostly involves an adjustment of the mitigation duration, rather than its magnitude, in response to a change in

32. As a result of this, figure 5 remains virtually unchanged if we assume testing is prohibitively costly in $t = 0$. We emphasize that this result must not be viewed as contradicting an expansive literature that underlines testing as a targeted mitigation policy to isolate only hazardous contacts. That is because testing in our model does not differ from other non-pharmaceutical interventions. This is different from saying that testing is an ineffective instrument when, *e.g.*, the disease has an asymptomatic phase, *à la* Berger et al. (2022).

initial infection.[33]

It is worth emphasizing that, when the initial infected population is equal to our estimation for Ontario in the early days of the pandemic in section 4.1, we find similar results to those in Glover et al. (2020): Calibrating a standard macroeconomic model within an epidemiological framework to the U.S. data, Glover et al. argue that

> a comparison of the utilitarian optimal policy to the actual policy in place as of Easter 2020 indicates that the shutdown in place was around twice as extensive as it should be. However, the optimal policy calls for leaving a partial shutdown in place well into the fall.

When we compare the optimal mitigation policy in figure 5 with the actual policy pursued in Ontario in the early days of the pandemic in figure 2, it appears that Ontario, too, reacted excessively to the pandemic's arrival. In addition, as in Glover et al., this initial reaction was followed by a premature reopening of the economy. These results change dramatically in the next section, under partial information.

4.3 Mitigation under Uncertainty

As the first step in our investigation of optimal policy under partial information, we consider a case where testing is prohibitively costly. This scenario is a reasonable description of the problem policymakers were facing in the starting months of COVID-19 pandemic, when COVID-19 testing kits were scarce, expensive and inaccurate, and health care systems had not yet developed a capacity to test people on a large scale.

Under the assumption of no testing ($\kappa = 0$), the positivity rate loses all its information content (since $\zeta(\kappa) \to \infty$ as $\kappa \to 0$). As a result, one can combine problems (15) with (16) as

$$
\begin{aligned}
V_t\left(N, H = h, Q^{beg} \sim \left(\mu^{beg}, \sigma^{beg}\right)\right) = & \\
\max_{m \in [0,1]} & \left\{ N \cdot u(Y/N) + \rho \mathbb{E}\left[V_{t+1}\left(N - \varphi(h), H', \left(Q^{beg}\right)' \right) \right] \right\} \\
s.t. \quad & Y = \Omega(m) \cdot w \cdot (N - h) - \Lambda(\kappa \cdot N).
\end{aligned}
\tag{24}
$$

33. Optimal policy converges in all cases after period 300, and demonstrate "waves" before the complete reopening of economy one month before the arrival of vaccine. This is because, as we get closer to the development of a cure, the "future costs" of the spread of the virus dissipate rapidly. This is consistent with Giannitsarou, Kissler, and Toxvaerd (2021)'s claim that optimal mitigation depicts oscillations before final reopening.

Figure 5. Time-Path of Policy, and Infected, Hospitalized, and Deceased Population under Optimal Policy

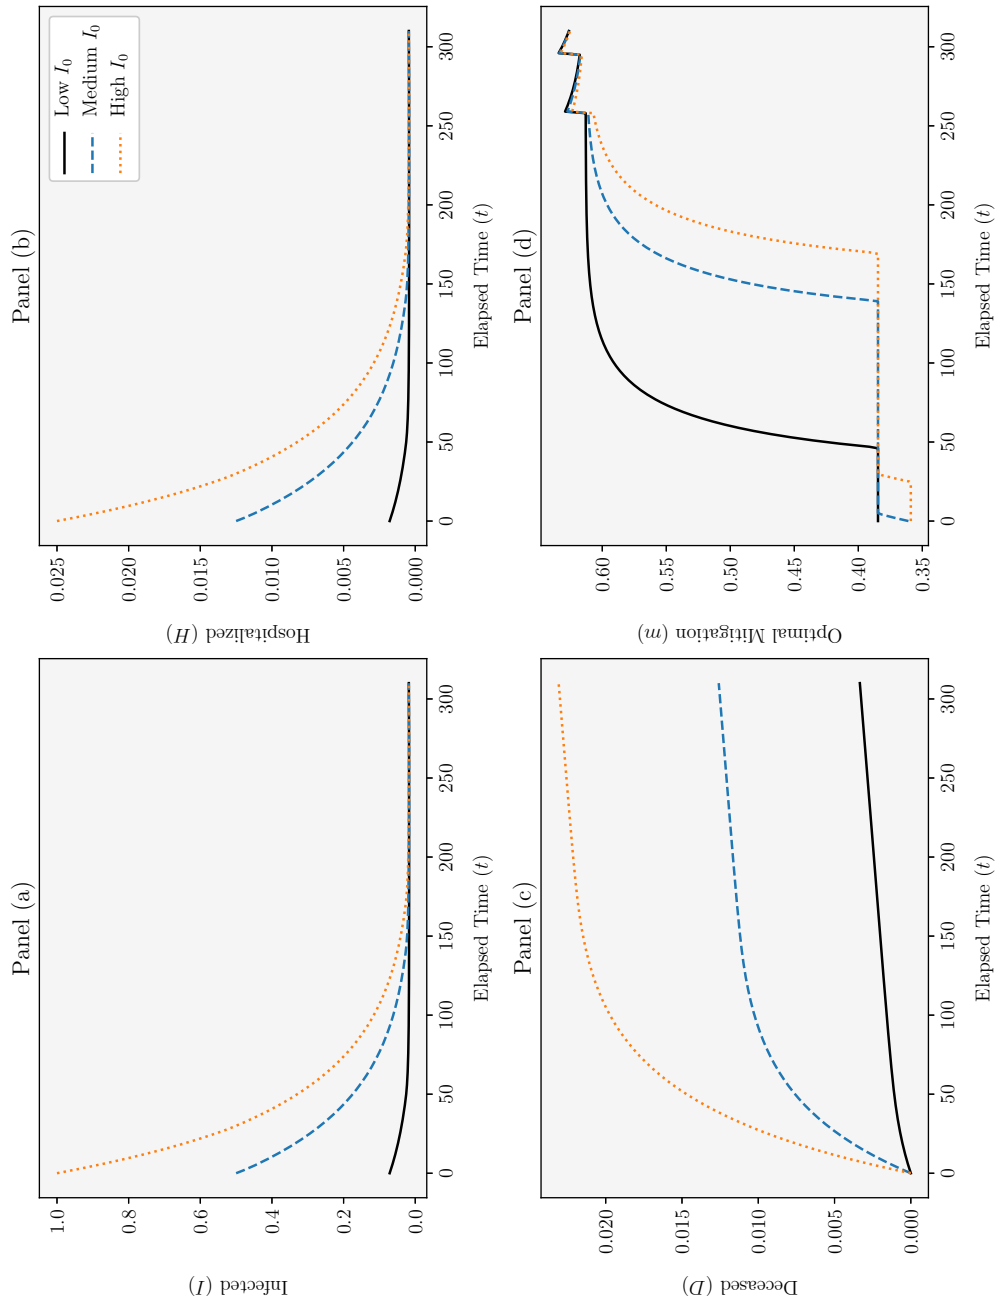

Note: Only the first 310 days are drawn. We assume planner has full information and that realized signal noise is zero in all periods. Low, medium and high I_0 correspond to 0.07, 0.5 and 1 percent of N_0, respectively. Discontinuities in optimal paths are, partly, artifacts of relatively coarse discrete grids and interpolation in the numerical exercise, and, partly, the result of the oscillatory nature of optimal mitigation. See the online appendix for an explanation.

In this problem, μ^{beg} serves as an indicator of the planner's beliefs about the infected. On the other hand, σ^{beg} captures the level of uncertainty in her beliefs. The evolution of μ^{beg} and σ^{beg} are specified by (10), (11) and (12) when $\kappa = 0$. In what follows, we show how different scenarios for these initial beliefs entail significantly different implications for the optimal mitigation policy.

Figure 6 illustrates a simulation exercise derived from solving problem (24), when the actual number of initially infected, I_0, is set equal to our estimate in section 4.1 for Ontario in the early weeks of the pandemic. While planner's beliefs coincide on average with I_0 ($\mu_0^{beg} = I_0$), there is moderate uncertainty about this number ($\sigma^{beg} = 0.1$ before observing H_0).

The left hand side axis in panel (a) in figure 6 depicts a sample path of *actual* infected population, together with the planner's mean beliefs about this variable. The right hand axis shows the number of hospitalized as time passes. As a result of the assumption that the *realized* noise in the signal is zero in all periods ($\varepsilon_t = 0$ for all $0 \leq t \leq T$), all the lines coincide in this graph, indicating that the planner is systematically making correct inferences about the underlying state.

As argued in section 3.2, the planner's choice of mitigation in panel (c), not only affects sample paths of I_t, μ_t^{beg} and H_t, but also the standard deviation of beliefs, as depicted in panel (b). In particular, when a strict lockdown is in place at the beginning, the level of uncertainty decreases dramatically. This is a direct result of the mitigation's role as an information acquisition device.[34]

A comparison of the optimal mitigation under partial information in figure 6 with that under full information (figure 5) highlights the potential role of uncertainty in the initial excessive reaction to the COVID-19 pandemic. The first two curves in figure 7 reproduce the optimal policy paths under these two scenarios. As the graph illustrates, uncertainty causes the planner to follow a significantly more strict mitigation policy—namely, an *optimal overreaction*, of up to 20% in the initial months.

This optimal overreaction, as one might expect, is a result of the role of mitigation in resolving uncertainty—as an information acquisition device. However, the planner's yearning for more information has a more subtle cause than an aversion to risk. To further elaborate this point, figure 8 compares the initial value functions for different levels of uncertainty, and for different realizations of H_0, as functions of average beliefs.

As figure 8 suggests, the planner's value function is convex in her *priors* about the average in-

34. A similar simulation in which hospitalization signals contradict planner's mean beliefs for two weeks shows that only a few periods of bad news is enough to permanently scar the planner's beliefs about the actual underlying state. This "scarring effect" decelerates the partial reopening after the initial 50 periods, when compared to the case without unfavourable signal noise.

Figure 6. Time-Path of Policy, and Actual Infected Population and State Variables under Optimal Policy

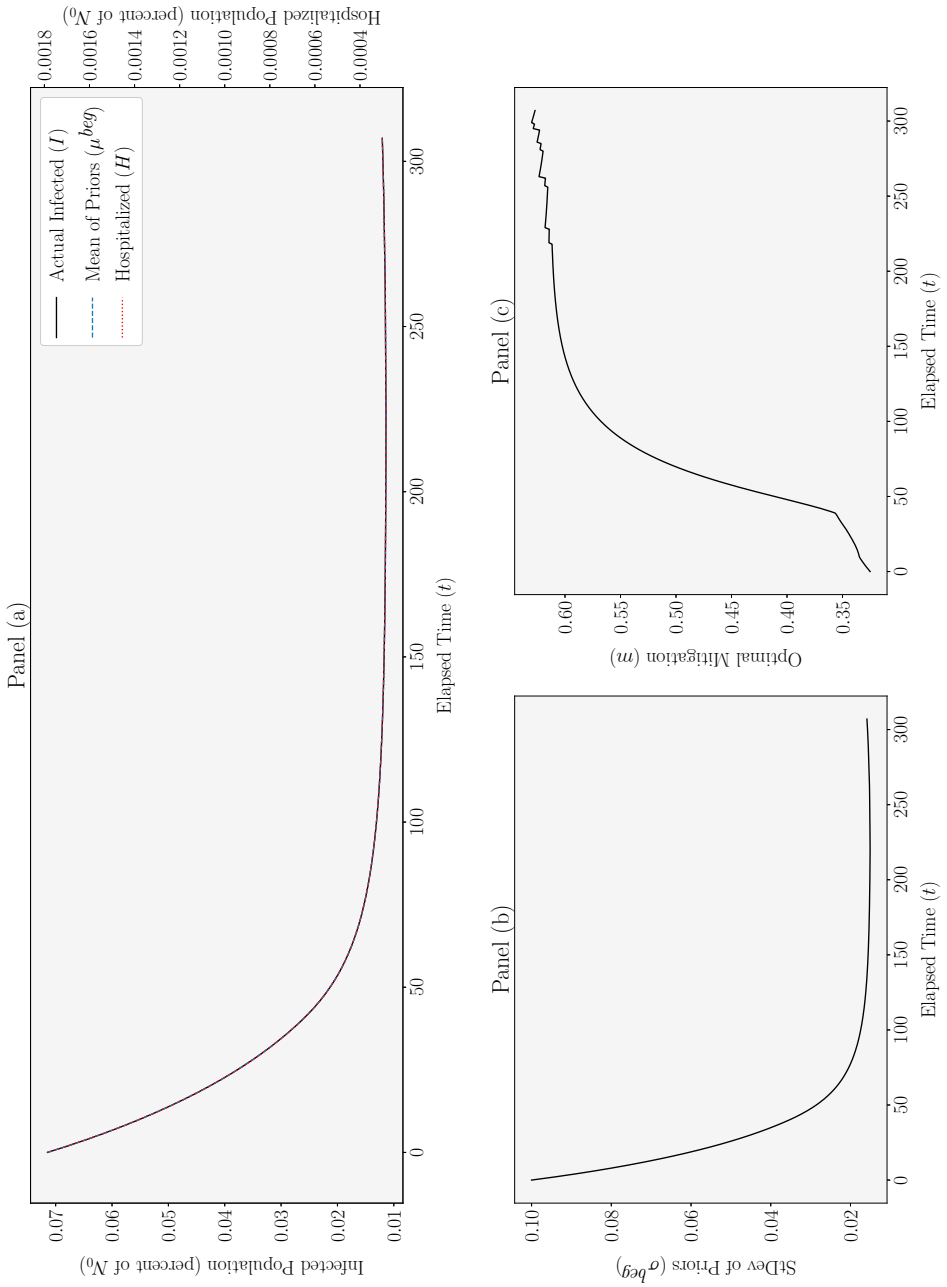

Note. We assume testing is prohibitively costly. Only the first 310 days are drawn. All curves in panel (a) coincide. Discontinuities in optimal paths are, partly, artifacts of relatively coarse discrete grids and interpolation in the numerical exercise, and, partly, the result of the oscillatory nature of optimal mitigation. See the online appendix for an explanation.

FIGURE 7. Time-Path of Optimal Policy under Different Scenarios

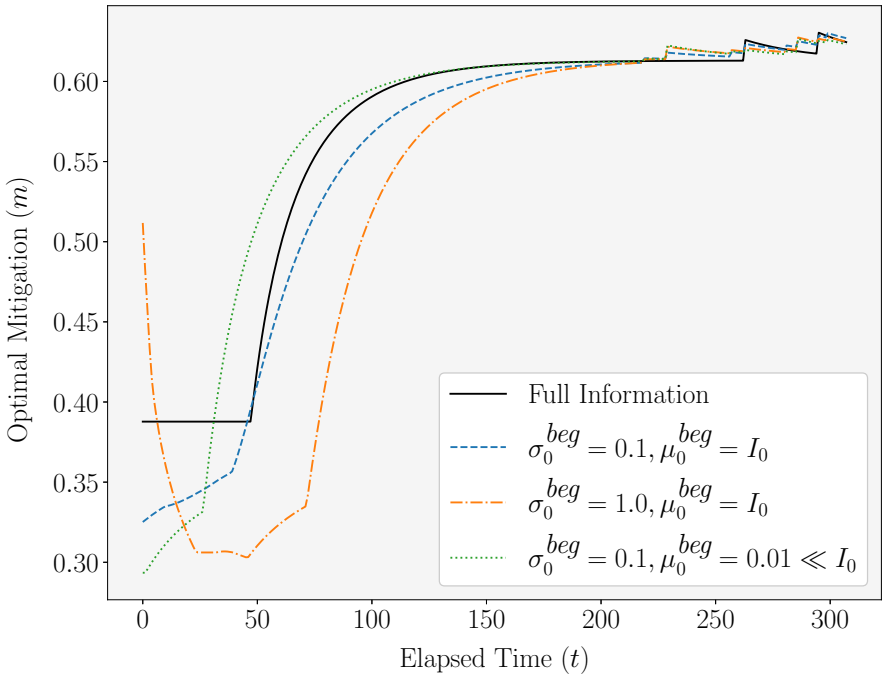

Note. Only the first 310 days are drawn. It is assumed that the testing is prohibitively costly, noise in the hospitalization signal is zero in all periods, and $I_0 = 0.07$ in all cases. Discontinuities in optimal paths are, partly, artifacts of relatively coarse discrete grids and interpolation in the numerical exercise, and, partly, the result of the oscillatory nature of optimal mitigation. See the online appendix for an explanation.

Figure 8. Value Functions under Different Levels of Uncertainty and Different Signals

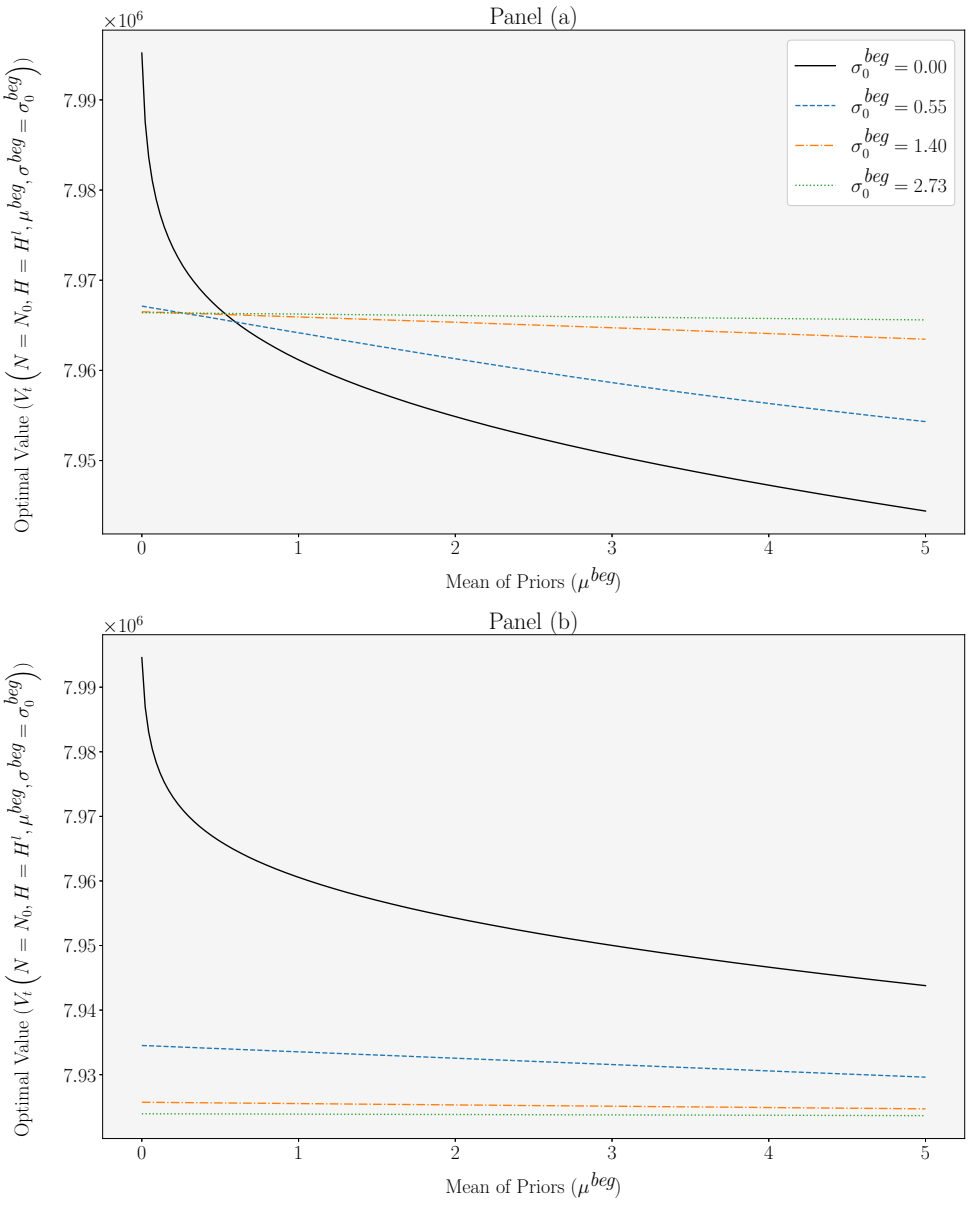

Note. Panels (a) and (b) correspond to $H_l = 0$ and $H^h = 0.36$, respectively, as percentages of N_0.

fected population, for a given realization of hospitalization.[35] As a result, the *ex ante* value is non-monotone in the variance of priors, for a given value of μ^{beg}. In addition, the shape of this dependence appears to change significantly once H_0 is realized. Theoretically, this is because of the fact that the posterior beliefs are a function of the prior's variance, and the shape of this function is highly dependent on the current realization of H_0.

To further expand on the intuition behind this observation, consider two planners: A planner who is certain that the underlying infected population is quite large, versus a planner who is rather uncertain about the underlying state. In case both of these planners observe a low level of hospitalization, the former attributes it to pure chance, while the latter uses it to infer that the underlying state is not as dire as once imagined. The second planner, consequently, has a higher expected utility than the former, be it due to her naivete.

The inverse of this argument also holds: For an uncertain planner who initially holds low beliefs about the underlying number of infected, the realization of a large number of people in the hospitals can be quite devastating. This is most clearly seen by comparing the value functions for a high value of σ^{beg}, in panels (a) and (b) of figure 8.

As figure 7 suggests, it appears that the sheer possibility of receiving an unfavourable hospitalization realization in the future is the dominant force in determining the planner's optimal course of action. This is in line with the intuition provided by the learning literature for the value of information in informing "at least as good decisions" as under uncertainty (Gollier 2001).

An experiment can highlight this force in the choice of policy further. The dashed-dotted orange line in figure 7 depicts the path of optimal mitigation under the same scenario as that in figure 6, but with considerable initial uncertainty ($\sigma^{beg} = 1.0$, as opposed to $\sigma^{beg} = 0.1$, in $t = 0$). As the figure suggests, optimal policy in this case involves a significant "under-reaction" in the first week or so. This initial lack of reaction, then, is compensated in the following weeks by subjecting the economy to severe restrictions which are then extended far beyond the case with only moderate uncertainty.

With significant initial uncertainty, mitigation loses its content as an information device, but still can generate costly policy mistakes. Therefore, the best course of action is to postpone the initial reaction by several periods until enough signals are received from the hospitals and the initial uncertainty is partially resolved.

As a final exercise, we consider a case where the initial uncertainty is accompanied by a consid-

35. This is different from the standard result in the leaning literature, that the payoff is convex in posterior beliefs, as shown by, *e.g.*, Gollier (2001).

erable underestimation of I_0, before observing the first signal ($\mu_0^{beg} = 0.01 \ll I = 0.07$). Optimal policy is shown by the dotted green curve in figure 7. Under this scenario, upon receiving the first indications that contradict her beliefs, the planner implements a strict mitigation policy. Compared to the case under full information, this policy exhibits a 35% overreaction.

Interestingly, as planner's beliefs converge to the actual number of infected, the policy responds by a premature reopening of the economy compared to the optimal policy under full information, similar to what Glover et al. (2020) document on policy response in the U.S. This is because of the rapid fall in the hospitalization signal as a result of the initial overreaction. This fall, coupled with the fall in uncertainty itself, inhibits planner's beliefs from further adjustments, leading to the relaxation of the restrictions sooner than what optimal policy suggests under full information.

Even though the magnitude of our results change to some extent as we change the parameters of the model—especially the value of being alive and the curvature of the utility function—the direction of these results are relatively robust to such changes. Our robustness tests seem to suggest that the rapid reopening of the economy, when μ^{beg} severely underestimate I_0, reverses for high values of ν. The magnitude of overreaction is also, at least up to some point, most sensitive to ν, even though the overreaction itself remains a robust feature of the constrained policy. Our results do not change in response to a change in the arrival time of the vaccine, T.[36]

4.4 Testing as an Information Instrument

Our results in the previous section signify the value of information to a planner under uncertainty, when policy mistakes can entail extreme costs. We showed that, when the only source of information to the planner is an exogenous noisy signal, the planner uses mitigation policy as an information acquisition device to avoid such policy mistakes.

These observations suggest that, had SARS-CoV-2 test kits been readily available at the start of this pandemic, large-scale random testing could have replaced overreaction in mitigation. In this section, we examine this hypothesis using our quantified model in a counterfactual scenario: We assume that testing is available to the planner at the onset of the pandemic at a cost.

Panels (a) and (b) in figure 9 show optimal paths of mitigation and testing, as solutions to problems (16) and (15), respectively. The initially infected population, average and standard devi-

36. One exception in our robustness tests is the value of I_0: Even though we estimate this value in our particle filtering exercise, after experimenting with high values of I_0, it appears that the overreaction disappears for levels of $\mu_0^{beg} = I_0$ in the order of 2 to 5 percent of the initial population. These values, however, are clearly outside any reasonable range.

ation of priors, and the realized noise in hospitalization signal are similar to those in figure 6 ($\mu_0^{beg} = I_0 = 0.07$, $\sigma_0^{beg} = 0.1$, and $\varepsilon_t = 0$ in all $t \geq 0$). In addition, we assume that the realized noise in the positivity rate is zero ($\upsilon_t = 0$ in all $t \geq 0$).

As one would expect, the convexity of the testing cost function limits the planner's utilization of this new instrument. Nevertheless, she still tests an astounding 14% of the entire population in the first period alone. To put this into perspective, such large-scale testing eats up more than 17% of the economy's output in a day, under the assumed cost function. This massive cost can be viewed as the *information premium* that the planner is willing to pay to avoid policy mistakes.

This relatively comprehensive testing policy leads to a large decline in the variance of *uncertainty*. Consequently, we would expect the mitigation policy not to exhibit much overreaction in this counterfactual economy. This is indeed the case after period $t = 0$.

What is surprising in figure 9 is the "under-reaction" in mitigation in period $t = 0$. Three factors contribute to this under-reaction: First, the comprehensive testing in $t = 0$ reduces the effective infection rate by about 12%. However, this is still far from locking down 60% of the economy in the case without uncertainty. Secondly, similar to the case with significant initial uncertainty, when initial uncertainty is large *relative to* the signal noise, the planner is willing to postpone mitigation to gain a more accurate insight into the underlying states. Finally, the cost that the planner is paying for testing in $t = 0$ is so large that the planner can no longer afford a lockdown. Nevertheless, information is so valuable to the planner that she is willing to pay such a large premium, allowing a liberal spread of the disease in $t = 0$.

5 Concluding Remarks

Across developed countries, initial reaction to the first confirmed cases of SARS-CoV-2 infection was strict and swift enactment of mitigation measures—such as lockdown of parts of or the entire economy. These initial policy responses were more excessive, in degree and range, when compared to the later stages of the pandemic, considering the fact later variants and waves demonstrated more concerning signs of a rapid and widespread propagation.

In this paper, we study the value of information during the pandemic in forming optimal mitigation and testing policies. We show that the above-mentioned excessive reaction can be justified through the lens of partial information about the true number of infected.

To that end, we develop an epidemiological model where the true number of infected cannot be directly observed, but inferred through two signals: hospitalization and positivity rate. An egalitarian planner can choose the mitigation measure and the scope of testing in order to control both

FIGURE 9. Time-Path of Optimal Policy in Presence of Testing

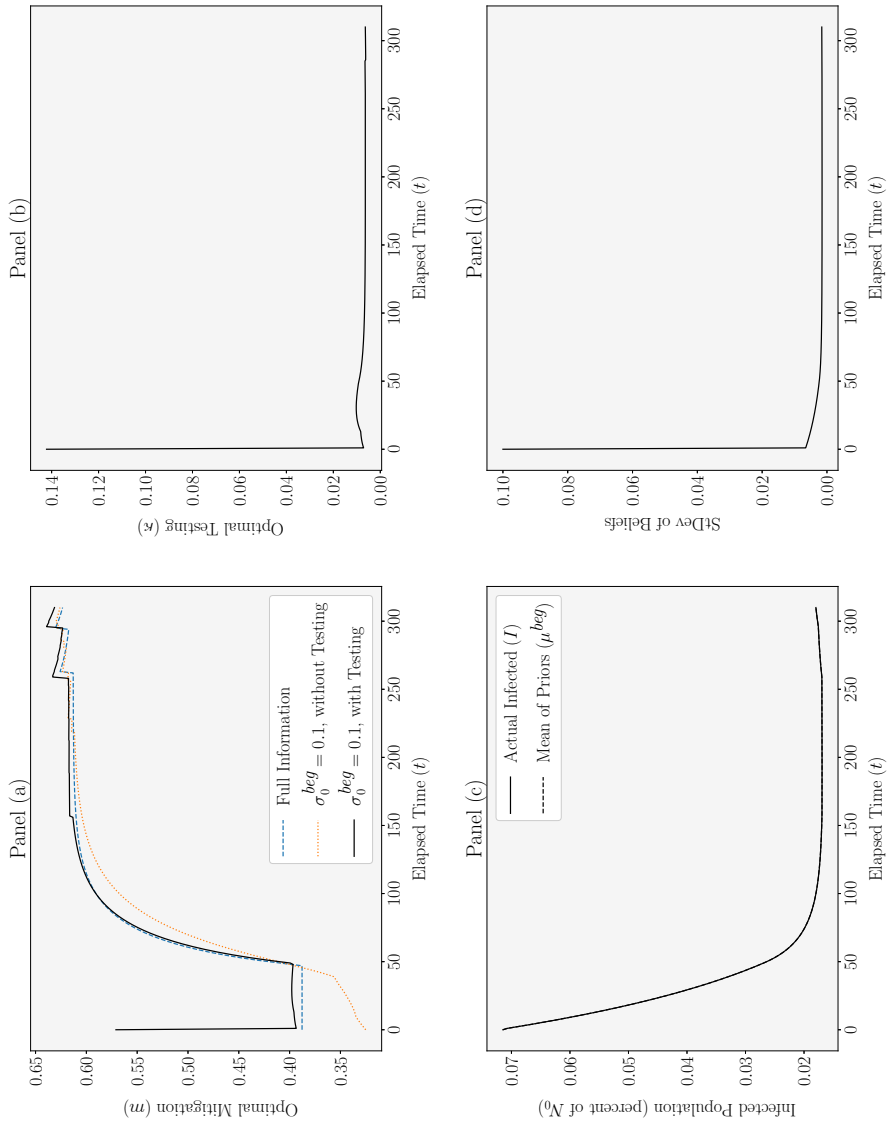

Note. Only the first 310 days are drawn. All black curves correspond to cases where testing is available from onset. Panel (a) also depicts optimal mitigation with full information and without testing for reference. All curves in panel (c) coincide. Realized noise in both signals are set to zero for all t. Discontinuities in optimal paths are, partly, artifacts of relatively coarse discrete grids and interpolation in the numerical exercise, and, partly, the result of the oscillatory nature of optimal mitigation. See the online appendix for an explanation.

the effective rate of infection and the quality of information regarding the number of infected.

Calibrated to the economy of Ontario, Canada, the model shows a similar degree of excessive reaction in mitigation—as compared to a counterfactual scenario without any uncertainty—when testing is prohibitively costly at the beginning of the pandemic and there is a reasonable degree of uncertainty about the number of infected; what we refer to as *optimal overreaction* in mitigation. We argue that this overreaction as the constrained optimum is the result of extreme costs of policy mistakes when there is uncertainty about the true number of infected. As such, overreaction in mitigation, serving as an information acquisition device in our environment, can prevent such policy mistakes.

In a counterfactual scenario, when testing is available at some cost from the beginning, we show that the planner is willing to allocate a significant portion of output for testing in order to eliminate the uncertainty. This is the *information premium* that the planner is willing to pay to know the true number of infected. Such premium replaces the information role of overreaction in our counterfactual, highlighting an overlooked role of testing as an information instrument during a pandemic.

Glover et al. (2020) argue that "one reason people disagree about the appropriate policy response [to the pandemic] is that the benefits and costs of lockdowns are large and very unequally distributed. Thus different groups prefer very different policies." By providing a general framework for the study of optimal public health policy under partial information during a pandemic, we provide an alternative justification for this: The optimal policy is extremely sensitive to our initial perceptions of the disease and to the costs our actions might entail for the society as a whole. As these beliefs and costs change, they shift the severity—or even the direction—of the optimal course of action.

At the end, we acknowledge that our model is still far from a realistic characterization of a policymaker's information constraints at the beginning of COVID-19 pandemic. Most importantly, our framework does not take into account that the planner might have been (and most probably was) uncertain about the underlying dynamics of the disease—its reproduction ratio, its true fatality rate, etc.[37] If we had factored in these considerations in our model, our results might have suggested a different initial response to the pandemic.

However, like most models of learning, even under simplifying assumptions, the resulting prob-

37. Let us not forget that even medical experts were not sure about the transmission mechanisms of SARS-CoV-2 one year into the pandemic. Had they advised on enacting mandatory face mask measures sooner, for instance, we could have avoided a natural disaster of such scales.

lem gets complicated quickly. This makes it quantitatively intractable to incorporate differing dimensions of partial information into the model. As a result, we decided to separate potential "structural uncertainties" from uncertainty about the state of the economy, and focus on the latter to underline the importance of testing to policymakers as an information device. Appendix F presents a simple framework for the study of structural uncertainty, when the underlying state is observed perfectly. The combined effect of these two dimensions of uncertainty and learning is an area worth exploring.

References

Ales, Laurence, Roozbeh Hosseini, and Larry E. Jones. 2014. *Is There "Too Much" Inequality in Health Spending Across Income Groups?,* National Bureau of Economic Research Working Paper.

Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang. 2020. "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic." *Journal of Public Economics* 191.

Alvarez, Fernando, David Argente, and Francesco Lippi. September 2021. "A Simple Planning Problem for COVID-19 Lock-down, Testing, and Tracing." *American Economic Review: Insights* 3, number 3 (): 367–82.

Angeletos, George-Marios, and Alessandro Pavan. 2009. "Policy with Dispersed Information." *Journal of the European Economic Association* 7 (1): 11–60.

Ash, Elliott, Sergio Galletta, Dominik Hangartner, Yotam Margalit, and Matteo Pinna. 2020. *The Effect of Fox News on Health Behavior during COVID-19,* Center for Law and Economics Working Paper Series 10/2020.

Atkeson, Andrew, Michael C. Droste, Michael Mina, and James H. Stock. October 2020. *Economic Benefits of COVID-19 Screening Tests,* NBER Working Papers 28031. National Bureau of Economic Research.

Bargain, Olivier, and Ulugbek Aminjonov. 2020. "Trust and Compliance to Public Health Policies in Times of COVID-19." *Journal of Public Economics* 192:104316.

Barrios, John M., Efraim Benmelech, Yael V. Hochberg, Paola Sapienza, and Luigi Zingales. 2021. "Civic Capital and Social Distancing during the Covid-19 Pandemic." *Journal of Public Economics* 193:104310.

Barrios, John M., and Yael Hochberg. 2020. *Risk Perception through the Lens of Politics in the Time of the COVID-19 Pandemic,* NBER Working Papers 27008. National Bureau of Economic Research.

Becker, Gary S., Tomas J. Philipson, and Rodrigo R. Soares. 2005. "The Quantity and Quality of Life and the Evolution of World Inequality." *American Economic Review* 95 (1).

Berger, David, Kyle Herkenhoff, Chengdai Huang, and Simon Mongey. 2022. "Testing and reopening in an SEIR model." *Review of Economic Dynamics* 43:1–21. ISSN: 1094-2025.

Brodeur, Abel, David Gray, Anik Islam, and Suraiya Bhuiyan. April 2021. "A Literature Review of the Economics of COVID-19." *Journal of Economic Surveys* ().

Brown, Robert Grover, and Patrick YC Hwang. 2012. *Introduction to Random Signals and Applied Kalman Filtering: with MATLAB Exercises.* Volume 4. John Wiley & Sons New York, NY, USA.

Bursztyn, Leonardo, Aakaash Rao, Christopher P. Roth, and David H. Yanagizawa-Drott. 2020. *Misinformation during a Pandemic,* NBER Working Papers 27417. National Bureau of Economic Research.

Chari, V. V., Rishabh Kirpalani, and Christopher Phelan. November 2020. "The Hammer and the Scalpel: On the Economics of Indiscriminate versus Targeted Isolation Policies during Pandemics." *Review of Economic Dynamics* ().

Cleevely, Matthew, Daniel Susskind, David Vines, Louis Vines, and Samuel Wills. 2020. "A Workable Strategy for COVID-19 Testing: Stratified Periodic Testing Rather than Universal Random Testing." *Oxford Review of Economic Policy* 36:S14–S37.

Data, MIT Election, and Science Lab. 2017. *U.S. President 1976–2020.* Harvard Dataverse.

Eberhardt, Jens Niklas, Nikolas Peter Breuckmann, and Christiane Sigrid Eberhardt. 2020. "Multi-Stage Group Testing Improves Efficiency of Large-Scale COVID-19 Screening." *Journal of Clinical Virology* 128:104382.

Einicke, Garry A., and Langford B. White. 1999. "Robust Extended Kalman Filtering." *IEEE Transactions on Signal Processing* 47 (9): 2596–2599.

Eslami, Keyvan, and Sayed M. Karimi. October 2019. *Health Spending: Necessity or Luxury? Evaluating Health Care Policies Using an Estimated Model of Health Production Function.* Working Paper.

Giannitsarou, Chryssi, Stephen Kissler, and Flavio Toxvaerd. 2021. "Waning Immunity and the Second Wave: Some Projections for SARS-CoV-2." *American Economic Review: Insights* 3 (3): 321–38.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull. 2020. *Health versus Wealth: On the Distributional Effects of Controlling a Pandemic.* Working Paper.

Gollier, Christian. 2001. *The Economics of Risk and Time.* MIT Press.

Gollier, Christian, and Olivier Gossner. 2020. *Group Testing against Covid-19,* EconPol Policy Brief 24. Leibniz Institute for Economic Research at the University of Munich.

Gonzalez, Francisco M., and Shouyong Shi. 2010. "An Equilibrium Theory of Learning, Search, and Wages." *Econometrica* 78 (2): 509–537.

Hall, Robert E., and Charles I. Jones. 2007. "The Value of Life and the Rise in Health Spending." *Quarterly Journal of Economics* 122 (1).

Huang, Guoquan P., Anastasios I. Mourikis, and Stergios I. Roumeliotis. 2008. "Analysis and Improvement of the Consistency of Extended Kalman Filter based SLAM." In *2008 IEEE International Conference on Robotics and Automation,* 473–479. IEEE.

Jones, Callum, Thomas Philippon, and Venky Venkateswaran. September 2021. "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home." *The Review of Financial Studies* 34, number 11 (): 5188–5223.

Jüni, P., A. Maltsev, G. M. Katz, A. Perkhun, and S. Yan. 2021. "Ontario dashboard." *Science Briefs of the Ontario COVID-19 Science Advisory Table.*

King, Aaron A., Dao Nguyen, and Edward L. Ionides. 2016. "Statistical Inference for Partially Observed Markov Processes via the R Package pomp." *Journal of Statistical Software, Articles* 69 (12): 1–43.

Kushner, Harold, and Paul G. Dupuis. 2014. *Numerical Methods for Stochastic Control Problems in Continuous Time.* Volume 24. Springer Science & Business Media.

Liptser, R. S., and A. N. Shiryaev. 2013. *Statistics of Random Processes I: General Theory.* Applications of Mathematics 5. Springer Science & Business Media.

Lucas Jr, Robert E. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4:103–124.

Melosi, Leonardo. 2017. "Signalling Effects of Monetary Policy." *Review of Economic Studies* 84 (2): 853–884.

Murray, Gregg R., and Susan M. Murray. 2020. "Following Doctors' Advice: Explaining the Issuance of Stay-at-Home Orders Related to the Coronavirus Disease 2019 (COVID-19) by US Governors." *OSF Preprints.*

Nimark, Kristoffer. 2008. "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics* 55 (2): 365–382.

Piguillem, Facundo, and Liyan Shi. April 2022. "Optimal Covid-19 Quarantine and Testing Policies." *The Economic Journal* 132, number 647 (): 2534–2562.

Schervish, Mark J. 2012. *Theory of Statistics.* 2nd edition. Springer Science & Business Media.

Simonov, Andrey, Szymon K. Sacher, Jean-Pierre H. Dubé, and Shirsho Biswas. 2020. *The Persuasive Effect of Fox News: Non-compliance with Social Distancing during the COVID-19 Pandemic,* NBER Working Papers 27237. National Bureau of Economic Research.

Subramanian, Rahul, Qixin He, and Mercedes Pascual. 2021. "Quantifying Asymptomatic Infection and Transmission of COVID-19 in New York City using Observed Cases, Serology, and Testing Capacity." *Proceedings of the National Academy of Sciences* 118 (9).

Townsend, Robert M. 1983. "Forecasting the Forecasts of Others." *Journal of Political Economy* 91 (4): 546–588.

Wells, Chad R., Jeffrey P. Townsend, Abhishek Pandey, Seyed M. Moghadas, Gary Krieger, Burton Singer, Robert H. McDonald, Meagan C. Fitzpatrick, and Alison P. Galvani. 2021. "Optimal COVID-19 Quarantine and Testing Strategies." *Nature Communications* 12 (1): 1–9.

Appendices

A Data

The data used in section 4.1 to estimate the epidemiological model using particle filtering is as follows:

1. COVID data

- Source: Government of Ontario
- Variables used in the paper

  - Number of patients hospitalized with COVID-19
  - Deaths
  - Confirmed Positive
  - Total patients approved for testing as of Reporting Date
- We normalize the number of hospitalizations and deaths with the total population (14,734,014) and fraction of population 30 and over (0.3562) sourced from Statistics Canada, 2020 Q3 population estimation.

2. Google mobility data

- Source: COVID-19 Community Mobility Reports
- Place id used in this paper: ChIJrxNRX7IFzkwRCR5iKVZC-HA
- Mobility index construction: A simple average of three categories
  - Retail and recreation
  - Workplace
- From the original mobility data, which is a percentage change compared to the baseline, we transform the data so that 1 is the baseline. ($\widetilde{x} = (x + 100)/100$ where $\widetilde{x}$ is the transformed data and x is the original data)

B Timing of Beliefs

As discussed in section 3.1, each period t in our economy can be divided into three stages according to the timing within a period: (i) the very beginning of a period, before the hospitalized population is realized; (ii) after the realization of hospitalization, before a testing decision is made; and (iii) after the positivity rate is realized, before a mitigation decision is made.

To see how the choice of stages affects the representation equation and, in turn, the number of states in our framework, let n identify the stage and period under consideration for the sake of discussion. Then, $\mathscr{I}^n$ is all the information available to the planner up to (and including) n (this includes the history of signals' realizations and policies), and $\mathscr{I}_n$ is the information that becomes available to the planner between stage n and the previous one (such that $\mathscr{I}^n = (\mathscr{I}^{n-1}, \mathscr{I}_n)$).

Let $Q_n(\boldsymbol{i}) \coloneqq \Pr(I_n \in \boldsymbol{i} \mid \mathscr{I}_n)$ be the probability of the infected population being in the set $\boldsymbol{i}$ at n, conditional on all the information available, capturing planner's beliefs about the underlying state of the economy.

In here and all that follows, when we talk about the planner's beliefs, we implicitly have a probability space in mind in which the sample space consists of the product of all possible sample paths of the observed and unobserved components of the underlying processes, a sigma-algebra over this set, and a probability measure over this space. The conditional probability measure, then, is a probability kernel over an appropriately defined filtration over this probability space.

For a rigorous treatment of this problem, see Liptser and Shiryaev (2013). For our purposes, it suffices to note that, under the assumptions of the model, and with a continuous initial belief, $Q(\cdot)$ remains absolutely continuous with respect to the Lebesgue measure and, as a result, has a Radon-Nikodym derivative—the so called probability density function of $Q(\cdot)$—which we will denote by small letter q, when of interest. The integrals with respect to the beliefs are understood to be Lebesgue integrals.

Then,

$$Q_{n+1}(\boldsymbol{i}) = \int \Pr(I_{n+1} \in \boldsymbol{i} \mid \mathscr{I}_{n+1}, I_n = i_n) \cdot Q_n(di_n). \tag{25}$$

Using the Bayes theorem, this updating equation can be written as

$$Q_{n+1}(\boldsymbol{i}) = A \cdot \int \Pr(\mathscr{I}_{n+1} \mid I_{n+1} \in \boldsymbol{i}) \cdot \Pr(I_{n+1} \in \boldsymbol{i} \mid I_n = i_n) \cdot Q_n(di_n), \tag{26}$$

for some constant of proportionality, A.[38] In what follows, we will refer to $Q_n(\cdot)$ and $Q_{n+1}(\cdot)$ as the planner's prior and posterior in n, respectively.

In general, the term $\Pr(I_{n+1} \in \boldsymbol{i} \mid I_n = i_n)$ in (26) is characterized by the *Fokker-Planck* equation. However, in our settings and with a deterministic law of motion of the infected population, this probability kernel boils down to a *Dirac delta* function for each value of i_n. It is clear that the point mass of this function depends on n and the choice of policy between n and $n+1$. An implication of this observation is that the choice of stages in which we keep track of the planner's beliefs have a critical bearing on the number of states in our problem formulation.

To economize on the set of states in our numerical analysis, thus, we pick stages (i) and (iii) to track the planning problem and beliefs.

38. See Schervish (2012) for a measure-theoretic formulation of Bayes rule.

C Particle Filtering

In section 4.1, we use a particle filtering method to estimate the following version of the epidemiological model:

$$I_t = (1-\gamma)I_{t-1} + \widetilde{\beta} m_{t-1} \left(\frac{S_{t-1}}{N_{t-1}}\right) I_{t-1} \tag{27}$$

$$S_t = S_{t-1} + \widetilde{\beta} m_{t-1} \left(\frac{S_{t-1}}{N_{t-1}}\right) I_{t-1} \tag{28}$$

$$H_t = \theta I_t + \varepsilon_t \tag{29}$$

$$K_t = \lambda I_t + \upsilon_t(\kappa_t) \tag{30}$$

$$N_t = N_{t-1} - D_{t-1} \tag{31}$$

$$\widetilde{\beta} = \beta(1 + \delta \mathbb{1}_{\{year>2020\}}) \tag{32}$$

$$\varepsilon_t \sim N(0, \nu) \tag{33}$$

$$\upsilon_t \sim \mathbb{N}\left(0, \zeta\left(\kappa_t\right)\right) \tag{34}$$

where D_{t-1} is the number of deaths due to COVID-19 and δ is the "variant infection factor" as explained in further detail in the Data and Calibration section (Section 4.1). Function $\zeta\left(\kappa_t\right)$ takes the form $\zeta\left(\kappa\right) = \kappa_{init} \cdot \nu_\kappa \cdot \left(\frac{1}{\kappa} - 1\right)$. where the standard deviation parameter ν_κ is normalized with $\kappa_{init} = 0.0408/100$, the (normalized) test rate on the initial data observation. All other variables follow the benchmark model equations.

The estimation employs a *method of iterated filtering* using `pomp` R-package, as follows:

- Pseudo code for the local and global search in Maximum Likelihood Estimation

  1. Set the initial guess $\Theta_0 = \{\nu_0, \nu_{\kappa 0}, \beta_0, \delta_0, I_{0,0}\}$ and prior distribution for the parameter set $\mathcal{M}(\Theta)$.
  2. Given Θ_i as the starting parameters, run the method of iterated filtering (`mif2`) N_l times **(local search)**.
  3. For N_l sets of resulting parameters from step 2., evaluate likelihoods of the process (27) using `pfilter`.
  4. From the prior distribution $\mathcal{M}(\Theta)$, randomly select N_g set of parameters.
  5. Using the N_g set of parameters from step 4. as starting points, run `mif2` N_g times **(global search)**.
  6. For N_g sets of resulting parameters from step 5., evaluate likelihoods using `pfilter`.

7. Find a new set of parameters Θ_{i+1} with the maximum likelihood among $N_l + N_g$ evaluations.
8. Stop if $i + 1 > N_{max}$. Otherwise go back to step 2.

- We use 3000 particles for each filtering and iterate the filtering process 300 times for each run of method of iterated filtering (`mif2`)
- We use $N_l = 20$, $N_g = 40$, and $N_{max} = 60$
- We use uniform distributions for all parameters as the prior.

$$\begin{aligned}
\nu &\sim \mathcal{U}(0.0001, 1) \\
\nu_\kappa &\sim \mathcal{U}(0.0001, 1.5) \\
\beta &\sim \mathcal{U}(0.01, 1) \\
\delta &\sim \mathcal{U}(0.001, 0.3) \\
I_0 &\sim \mathcal{U}(0.001, 30)
\end{aligned}$$

- Log likelihood of is evaluated using a sequential Monte Carlo algorithm (`pfilter` in `pomp`)
- For more details of the pomp, see King, Nguyen, and Ionides (2016).

D SIR vs SIS Models

Our choice of epidemiological model in section 3.1 is a slight variation of a SIS model, captured by the following system of equations

SIS Model

$$I' = (1 + \beta - \gamma) I, \tag{35a}$$
$$H = \theta I + \varepsilon, \tag{35b}$$
$$N' = N - \varphi(H), \tag{35c}$$

absent any policy intervention. This is different from the standard compartmental model used by most epidemiologists to model the dynamics of COVID-19 pandemic, the so-called SIR model. In

the context of our economy, the SIR model is characterized by the following system:

SIR Model

$$I' = (1-\gamma)\, I + \beta \left(\frac{IS}{N}\right), \tag{36a}$$

$$S' = S - \beta \left(\frac{IS}{N}\right) + \xi R, \tag{36b}$$

$$R' = (1-\xi)\, R + [\gamma I - \varphi(H)], \tag{36c}$$

$$H = \theta I + \varepsilon, \tag{36d}$$

$$N' = N - \varphi(H). \tag{36e}$$

In these equations, R is the population that fully recovers from the disease and enters into the economy again. In this variation of the SIR model, parameter ξ captures the rate at which the recovered population becomes susceptible to the disease again (which appears to be a characteristic of SARS-CoV-2, as we note in footnote 9).

As discussed at the end of section 3.1 and, in detail, in section 3.2, the linearity of the law of motion of infected population in (35a) is a critical condition for the tractability of the representation equation. Otherwise, we cannot capture the planner's beliefs by a well-behaved distribution any longer, and the planner's problem becomes numerically intractable as a result.

One standard approach in the optimal filtering literature to deal with such intractabilities—*e.g.*, when dealing with a non-linear equation like (36a)—is to approximate the law of motion of state by a linear one (and noise involved by a normal one). Another possibility is to approximate the experimental beliefs by a Gaussian distribution in each period. In this paper, we take the first approach, which forms the idea behind the extended Kalman filters.

Beside its technical necessity, one can argue that, in the period under consideration (that is the starting months of the pandemic when uncertainty is at its peak), the SIS system provides a good approximation of the SIR: If (i) the number of infected population does not exceed a threshold, and (ii) the recovered population becomes susceptible in the near future, S_t/N_t remains reasonably close to unity. Consequently, one state variable drops out of the control system, and, more importantly, equation (36a) can be approximated by a linear equation like that in (35a). The latter plays a critical role in keeping the planner's beliefs tractable throughout our analysis, as discussed before.

In this section, we demonstrate that, under our parameters estimates, our "linearization" remains an accurate one around the SIR model in the first three months of the COVID-19 pandemic, even in the absence of mitigation. Given that our quantitative results focus on the initial months of

the pandemic, we believe that our approximation remains reasonably accurate in addressing the questions that are of interest to this study.

Figure 10 illustrates the evolution of infected population, resulting from simulating the two systems under our estimated parameter values of section 4.1, when $\xi = 0$, $N_0 = 100$, $I_0 = 0.1$ and absent any mitigation measure. As one expects, and as panel (a) of the graph confirms, the SIS model predicts a divergence of infected population, whereas, "herd immunity" is achieved in the SIR model when I reaches a certain threshold. However, as panel (b) of figure 10 suggests, for the initial three months of the pandemic, the SIS model provides a rather close approximation to the SIR model.

We can expect this approximation to improve further with the introduction of mitigation measures (reducing the effective infection rate) and the possibility of reinfection ($\xi > 0$). Figure 11 demonstrates this by comparing the results of figure 10 to that when it takes 90 days, on average, for a recovered individual to become susceptible to the disease again.

When restrictive mitigation policies are in place, the SIS model becomes quite a reasonable approximation to the SIR model, not just at the beginning of the pandemic. This point is demonstrated in figure 12, where we have assumed a 30% lockdown policy in effect. As our discussions in section 2 suggest, this is a lower bound for the fall in mobility, *e.g.*, in Ontario, Canada. Even for this conservative choice of mitigation policy, the SIS model remains a reliable approximation to the SIR model until a year into the pandemic. As one can imagine, for stricter lockdown policies in effect—*e.g.*, around the magnitudes mandated by an overreacting planner—the two models converge.

At the end, it is worth emphasizing that such approximations are crucial and common in solving optimal filtering problems. For instance, as Kushner and Dupuis (2014) explain, the linearization introduced above uses the fundamental idea behind the *Kalman filter*, a rare occasion in which the representation equation becomes tractable, and has been used extensively before in *extended Kalman filters*.[39]

E An Alternative Timing

As discussed in footnote 18, an alternative to the timing considered in section 3.1 is to assume the planner has to decide about κ and m simultaneously, and before observing the positivity rate, K.

39. See Einicke and White (1999), Huang, Mourikis, and Roumeliotis (2008), and Brown and Hwang (2012) for discussions of extended Kalman filters and their properties.

Figure 10. SIR vs SIS Models without Reinfection and Mitigation

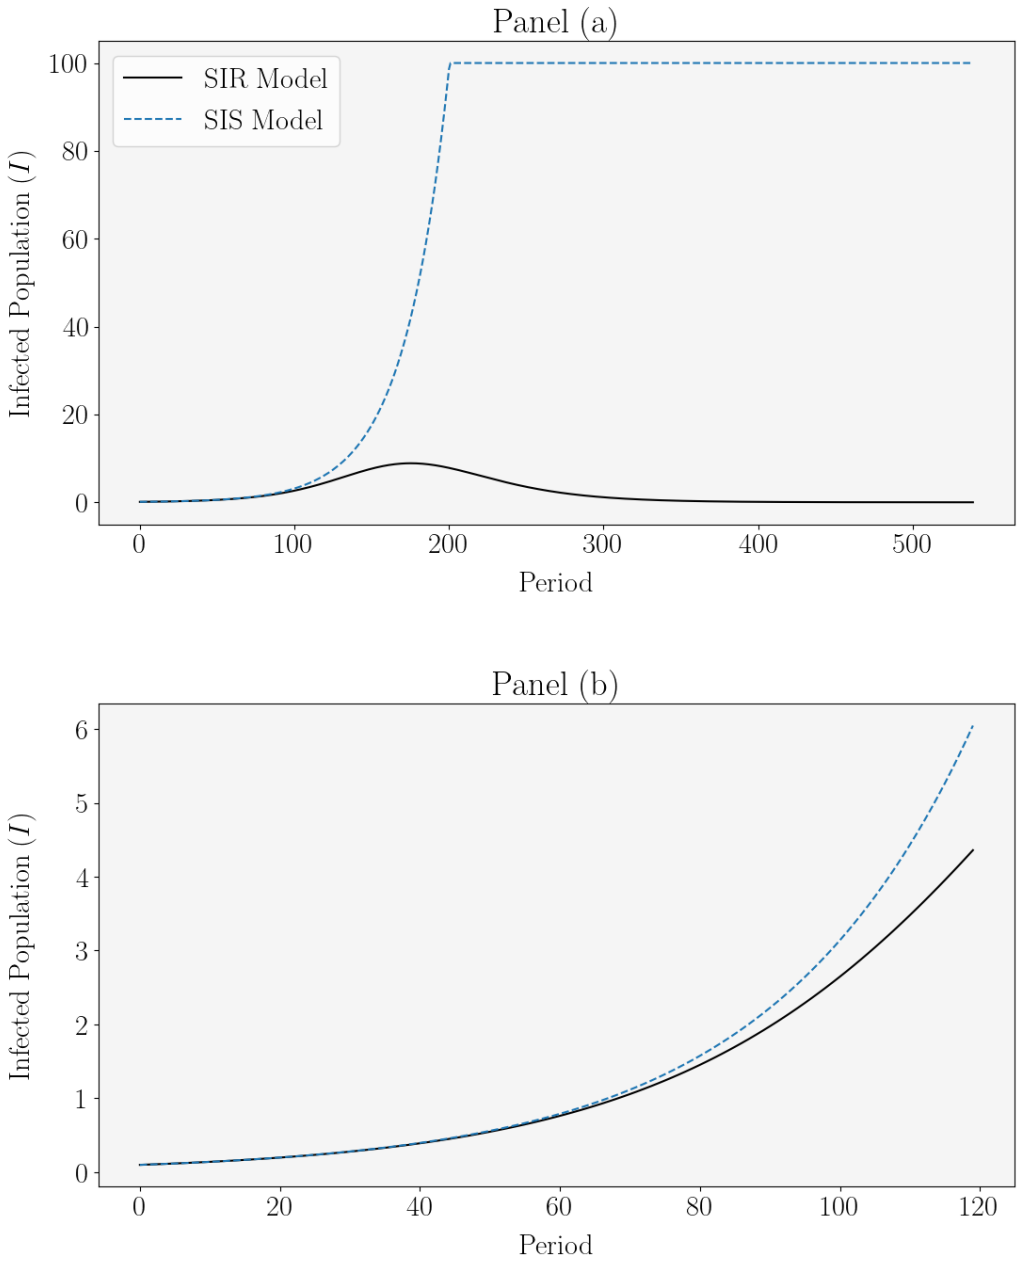

Note. It is assumed that $N_0 = 100$ and $I_0 = 0.1$.

Figure 11. SIR vs SIS Models with Reinfection

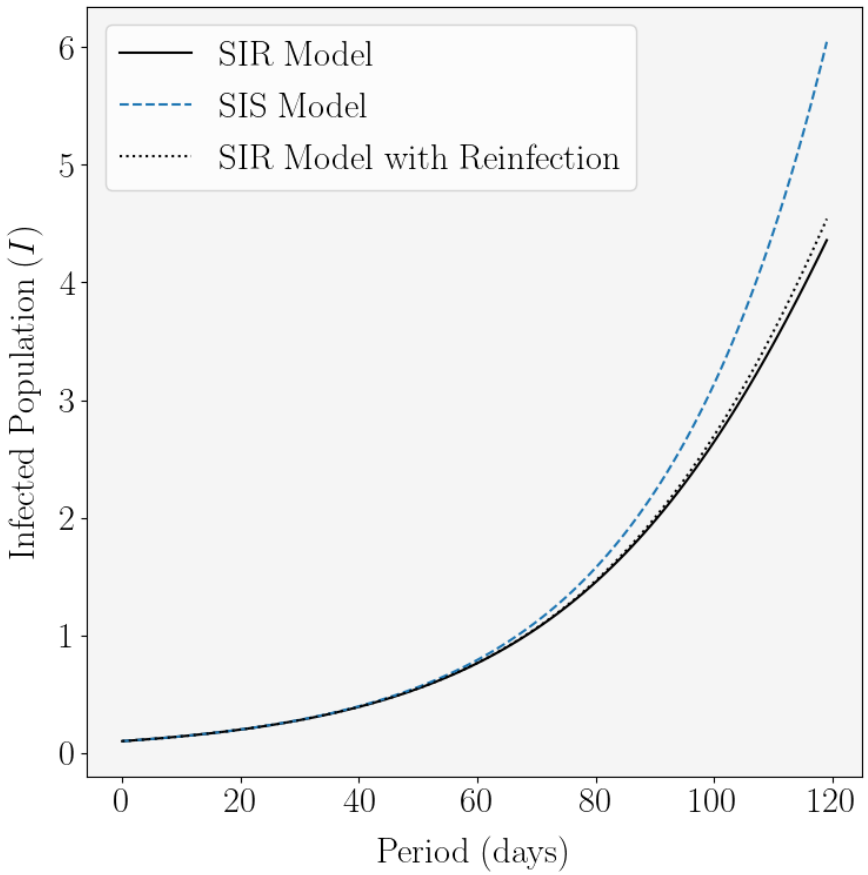

Note. It is assumed that $N_0 = 100$, $I_0 = 0.1$ and $\xi = 1/90$.

Figure 12. SIR vs SIS Models with Mitigation

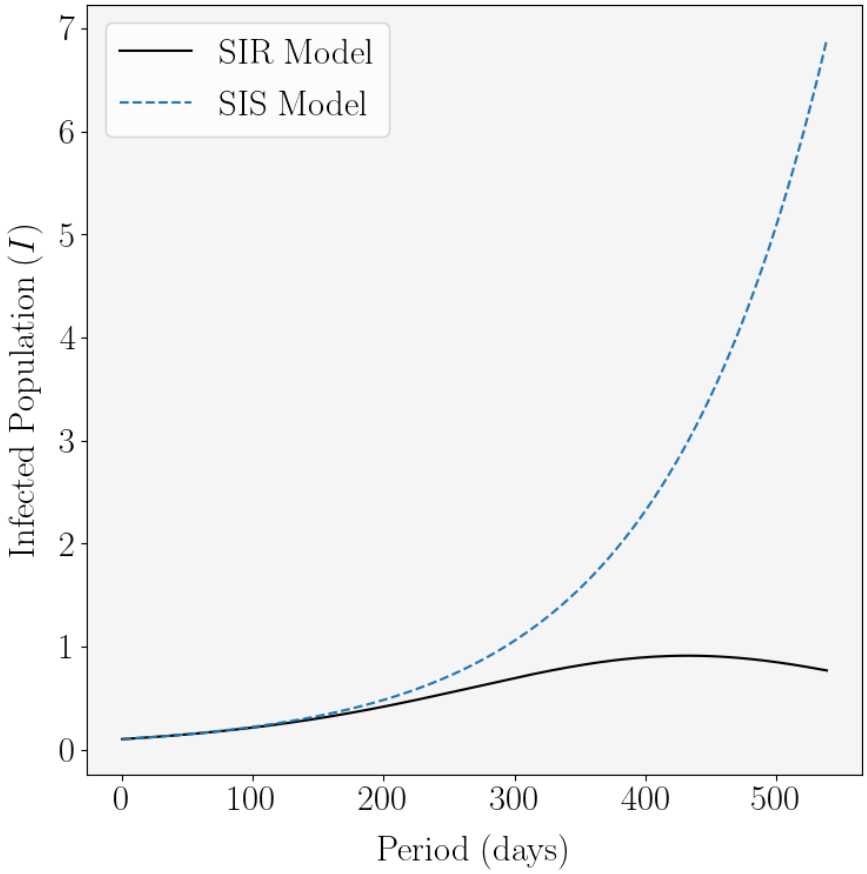

Note. It is assumed that $N_0 = 100$, $I_0 = 0.1$ and $\widehat{\beta} = 0.7\beta$.

Under this alternative timing assumption, the planning problems in (15) and (16) need to be combined together as a single problem at the beginning of period as follows:

$$\begin{aligned} V_t\left(N, H=h, Q^{beg} \sim \mathbb{N}\left(\mu^{beg}, \sigma^{beg}\right)\right) \\ = \max_{(\kappa,m)\in[0,1]^2} \Big\{ & Nu\left(\left[\Omega\left(m\right) w\left(N-h\right) - \Lambda\left(\kappa N\right)\right]/N\right) \\ & + \rho\mathbb{E}\left[V_{t+1}\left(N-\varphi\left(h\right), H', \left(Q^{beg}\right)' \sim \mathbb{N}\left(\left(\mu^{beg}\right)'\left(K\right), \left(\sigma^{beg}\right)'\right)\right)\right] \Big\}, \end{aligned} \tag{37}$$

subject to

$$\left(\sigma^{beg}\right)' = \left[1 + m\Gamma\left(\kappa\right)\beta - \gamma\right] \left[\frac{\zeta\left(\kappa\right)\left(\frac{\nu\sigma^{beg}}{\sqrt{\theta^2\left(\sigma^{beg}\right)^2+\nu^2}}\right)}{\sqrt{\lambda^2\left(\frac{\nu^2\left(\sigma^{beg}\right)^2}{\theta^2\left(\sigma^{beg}\right)^2+\nu^2}\right) + \zeta\left(\kappa\right)^2}}\right], \tag{38}$$

$$\left(\mu^{beg}\right)'\left(K\right) = \left[1 + m\Gamma\left(\kappa\right)\beta - \gamma\right] \left[\frac{\lambda k\left(\frac{\nu^2\left(\sigma^{beg}\right)^2}{\theta^2\left(\sigma^{beg}\right)^2+\nu^2}\right) + \left(\frac{\theta h\left(\sigma^{beg}\right)^2+\mu^{beg}\nu^2}{\theta^2\left(\sigma^{beg}\right)^2+\nu^2}\right)\zeta\left(\kappa\right)^2}{\lambda^2\left(\frac{\nu^2\left(\sigma^{beg}\right)^2}{\theta^2\left(\sigma^{beg}\right)^2+\nu^2}\right) + \zeta\left(\kappa\right)^2}\right], \tag{39}$$

and the terminal condition

$$V_T\left(N, H=h, Q^{beg} \sim \mathbb{N}\left(\mu^{beg}, \sigma^{beg}\right)\right) = N\left[u\left(\left(\frac{N-h}{N}\right)w\right) + \rho\frac{u\left(w\right)}{\left(1-\rho\right)}\right], \quad \forall Q^{beg}. \tag{40}$$

The major difference between problem (37) and the planning problem in section 3.1 is that, now, the expectation is with respect to the joint distribution of K and H'. Noting that K and H' are independent random variables, we can use the *law of iterated expectations* to compute this expectation in two stages: First, with respect to K, and, next, with respect to H' for each given value of $\left(\mu^{beg}\right)'\left(K\right)$ and $\left(\sigma^{beg}\right)'$.

Formally, we can write the expectations in (37) as

$$\begin{aligned} \mathbb{E}\left[V_{t+1}\left(N-\varphi\left(h\right), H', \left(Q^{beg}\right)' \sim \mathbb{N}\left(\left(\mu^{beg}\right)'\left(K\right), \left(\sigma^{beg}\right)'\right)\right)\right] \\ = \mathbb{E}_K\left[\mathbb{E}_{H'|K}\left[V_{t+1}\left(N-\varphi\left(h\right), H', \left(Q^{beg}\right)' \sim \mathbb{N}\left(\left(\mu^{beg}\right)'\left(K\right), \left(\sigma^{beg}\right)'\right)\right)\right]\right]. \end{aligned} \tag{41}$$

Using Bayes' rule, the two probability distributions in (41) can be derived as

$$\Pr\left(K=k\right)=\left[\frac{1}{2\sqrt{\pi\sqrt{\left(\sigma^{int}\right)^{2}\lambda^{2}+\zeta\left(\kappa\right)^{2}}}}\right]\cdot\exp\left(-\frac{1}{2}\left[\frac{k-\lambda\mu^{int}}{\sqrt{\left(\sigma^{int}\right)^{2}\lambda^{2}+\zeta\left(\kappa\right)^{2}}}\right]^{2}\right). \tag{42}$$

and

$$\Pr\left(H'=h'\mid K=k\right)=\left[\frac{1}{2\sqrt{\pi\sqrt{\left(\left(\sigma^{beg}\right)'\right)^{2}\theta^{2}+\nu^{2}}}}\right]\times\exp\left(-\frac{1}{2}\left[\frac{h'-\theta\left(\mu^{beg}\right)'\left(k\right)}{\sqrt{\left(\left(\sigma^{beg}\right)'\right)^{2}\theta^{2}+\nu^{2}}}\right]^{2}\right). \tag{43}$$

Solving problem (37) follows a similar procedure as the one used in solving (15) and (16). From a numerical standpoint, however, this problem is easier to solve than the latter two from one aspect, and harder, from another: First, since κ does not enter into (37) as a state variable and since the planning problem involves only one stage, the current timing can be easier to handle. On the other hand, the optimization step in problem (37) involves going over a two dimensional grid. This makes the optimization step rather time-consuming.

It turns out that having one fewer state variables and one stage (instead of two) deems the current timing easier to solve numerically than the one in section 3.1, for small values of T. Our quantitative results remain more or less the same.

However, as T gets large and/or the control grid becomes finer, the fact that optimization is on a 2D grid in this timing turns out detrimental to drawing reliable conclusions about the results for the following reason: In this case, we take advantage of the graphical processing unit (GPU)'s power to find the optimal policy in each step—using CUDA programming language. This is necessary to be able to solve such a large-scale program with such a massive control grid.

As the control grid becomes finer, we no longer have enough internal memory in the GPU unit to keep the policy function, in its entirety, in the unit. Our novel trick, in our benchmark timing, is to compress and then offload the policy functions from the RAM to an external storage unit, before moving to the next iteration. This, however, turns out to be extremely time consuming when one wants to offload the policy from the GPU's internal memory.

As T increases, something similar occurs: This time, the size of the external file holding the policy

function gets out of hand, even at the maximum level of compression.

As such, our results in this section must be taken with some level of caution. In addition, we tend to think of our timing in section 3.1 as a more reasonable description of what policymakers were facing around the world during the COVID-19 pandemic.

F Structural Uncertainty

Our environment—as laid out in section 3—makes a strong assumption about the nature of uncertainty: That the planner is uninformed about the "state" of the economy and the disease, while is certain about the "structure" of both.

While this assumption enables us to focus on the role of testing as an information instrument, it abstracts from the reality that most policymakers were quite uncertain about the nature of the SARS-CoV-2 virus at the onset of the pandemic. This, above all, included uncertainty about the infection rate which was much debated (at least) up to the point when COVID-19 was declared a global health crisis a few months after the first cases were diagnosed.

The advantage of this simplification is that, without it, our model becomes numerically intractable given our available computational resources. On the downside, neglecting the *structural uncertainty* can well explain part of the policy reaction to the pandemic at its onset, either reinforcing or negating our results in section 4.

To investigate this possibility, in this appendix, we present a version of the model in which uncertainty is about the infection rate, β, rather than the state of the economy. The planner has to form and adjust her beliefs about the true infection rate based on the observed evolution of the infected population in the economy which is, now, contaminated by noise. In the next subsection, we investigate the implications of this structural uncertainty for optimal policy

Consider the same baseline environment as in section 3, but, now, assume that β can be an element of the set $\mathcal{B} \coloneqq \{\beta_L, \beta_H\}$, where $\beta_L < \beta_H$. To capture the planner's beliefs, we denote her perceived chances of $\beta = \beta_H$ by π.

Unlike our baseline economy, the planner perfectly observes the infected population I_t in each period t. As such, testing no longer plays an informational role in our economy, and is consolidated by other mitigation measures as a single instrument which we keep denoting by m. We assume, as before, that $m \in [0, 1]$, where $m = 0$ indicates a full shutdown of the economy. We capture the cost of consolidated mitigation by function $\Omega(\cdot)$, with a slight abuse of notation.

Given the true value of β and the mitigation policy m, the law of motion of infection is given as

$$I_t = (1 + m \cdot \beta - \gamma) \cdot I_{t-1} + \varepsilon_t, \tag{44}$$

where ε_t is an independently and identically distributed noise with mean zero, whose distribution is captured by $G(\cdot)$

Here, we assume that fraction θ of the infected population ends up in hospitals in each period with certainty. Of this population, some succumb to the disease by the end of the period. This fatality function is given by

$$\varphi(I) = \min\left\{\varphi_1 \cdot (\theta I) + \varphi_2 \cdot (\theta I)^2, \theta I\right\}. \tag{45}$$

We assume the remainder of the economy to be similar to that of section 3. The timing of the problem is as follows: At the beginning of period t, the planner observes the infected population I_t, and updates her beliefs accordingly. Next, she makes a mitigation decision, m_t, based on which production and consumption take place. At the end of the period, $\varphi(I_t)$ of the infected die and leave the economy forever.

The Planner's Problem

As in our baseline model, the choice of the stage within a period in which we choose to write the planning problem has a bearing on the number of state variables in the planning problem. To economize on the number of states, we consider the planning problem after the realization of I_t and after the beliefs, π, have been updated accordingly.

This problem can be written as

$$\begin{aligned} V_t(N_t, \pi_t, I_t) = \max_{m_t \in [0,1]} \quad & \left\{ N_t \cdot u(Y/N_t) + \rho \cdot \mathbb{E}\left[V_{t+1}\left(N_t - \varphi(I_t), \pi_{t+1}(I_{t+1}), I_{t+1}\right)\right] \right\} \\ s.t. \quad & Y = \Omega(m_t) \cdot w \cdot (N_t - \theta I_t). \end{aligned} \tag{46}$$

The expectation in problem (46) is with respect to the number of infected in $t+1$. Since π in this formulation represents beliefs after observing the infected population, it is written as a function of I_{t+1} to emphasize its dependence on the received signal. This dependence is determined by Bayes' theorem as

$$\pi_{t+1}(I_{t+1}) = \frac{\pi_t \cdot \Pr(I_{t+1} = i \mid \beta = \beta_H, I_t)}{(1 - \pi_t) \cdot \Pr(I_{t+1} = i \mid \beta = \beta_L, I_t) + \pi_t \cdot \Pr(I_{t+1} = i \mid \beta = \beta_H, I_t)}. \tag{47}$$

FIGURE 13. Initial Value as Function of Infected Population

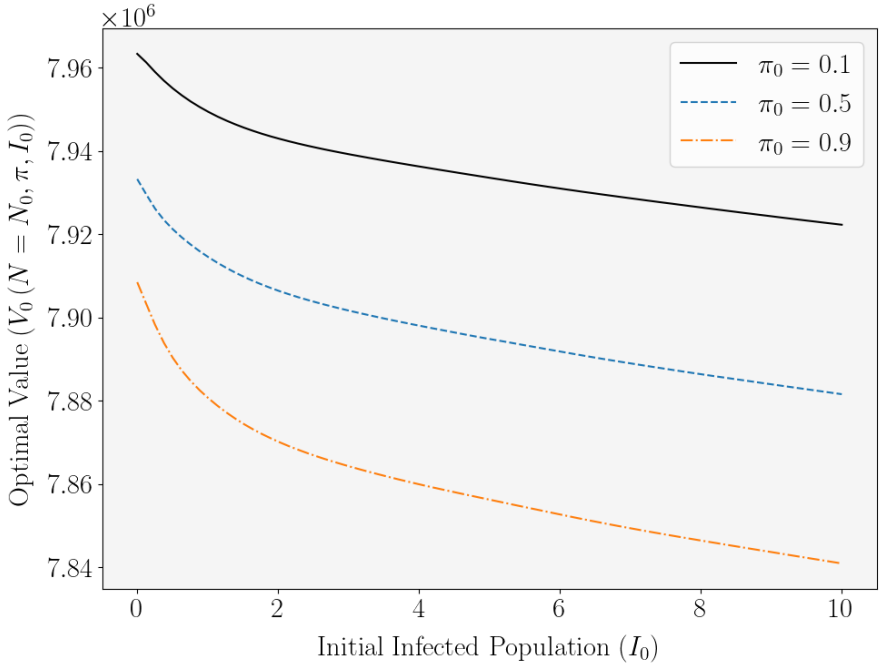

A Numerical Example

As in the baseline model, finding a closed-form solution to (46) is not possible. Therefore, in this section, we discuss the properties of the solution via a numerical example, similar to that in section 4.

To this end, we use the same parameter values for the utility function, discount rate, fatality function and openness function as those in section 4. We assume $G(\cdot) \sim \mathbb{N}(0, \nu)$. We assume a baseline value of ν of 0.03, but also check the robustness of our results for $\nu = 0.1$ and $\nu = 0.3$. We set $\beta_L = 0.091$ and $\beta_H = 0.191$, and choose γ, θ, w and T before.[40]

Figure 13 illustrates the initial value function, $V_0\,(N = N_0, \pi, I_0)$, for three different values of π_0, as functions of initial infected population, I_0. One must note that, unlike the value function in our baseline model, I_0 is observed perfectly, and π_0 represents the probability that the planner assigns to $\beta = \beta_H$ *after* the realization of I_0.

As the figure suggests, the value demonstrates the same convexity in the number of infected as

40. Note that, unlike our parametrization in section sec:quant, these parameter values are not estimated, but picked intuitively for the sake of demonstration.

that in the baseline model. As such, like before, everything else the same, the planner prefers a mean-preserving spread over the infected population over a certain outcome. In the current setting, however, uncertainty about the future realizations of I stem from two sources: Structural uncertainty about the underlying infection rate (captured by π), and the signal noise (as captured by ν).

As such, mitigation in this setting has two opposing effects on the future uncertainty. First, depending on the relationship between β_L and β_H, and the distribution of noise, ε, less mitigation (m closer to one) can help eliminate uncertainty faster than more strict mitigation measures (m closer to zero).

To see this most clearly, suppose $\beta_L \ll \beta_H$. Then, $m = 1$ can effectively separate β_L from β_H in a single period, if the variance of ε—implied by $G(\cdot)$—is so small that

$$\Pr\left(\overline{I} < I' \mid I, \beta = \beta_L, m = 1\right), \Pr\left(I' < \overline{I} \mid I, \beta = \beta_H, m = 1\right) < \upsilon,$$

for some arbitrarily small $\upsilon > 0$. On the other hand, if $m = 0$, then

$$\Pr\left(I' = \overline{I} \mid I, \beta = \beta_L, m = 0\right) = \Pr\left(I' = \overline{I} \mid I, \beta = \beta_H, m = 0\right),$$

under the same assumptions.

This can also be seen in the updating equation, equation (47): When either $\Pr(i' \mid \beta = \beta_L, I)$ or $\Pr(i' \mid \beta = \beta_H, I)$ is zero, then π' is either equal to one or zero. (Note that $\pi = 1/2$ implies the highest uncertainty in this environment.) Therefore, less mitigation can quickly resolve uncertainty by pushing π' towards either zero or one.

However, less mitigation implies rather large uncertainty about the infected population in the immediate future. To see this, suppose π is close to $1/2$. Then, with $m = 1$, the distribution of I' becomes a *bimodal distribution*—an equally weighted sum of two normal distributions with the same variances but distinct means. With $m = 0$, however, $I' \mid I = i \sim \mathbb{N}((1-\gamma)i, \nu)$.

As such, in choosing the optimal mitigation, the planner has to balance these two opposite incentives, against the standard trade-off that exists between the cost of allowing the virus to propagate versus the cost of mitigation. (One must note that the preceding arguments are under the assumptions of a rather small ν and $\beta_L \ll \beta_H$.)

Figures 14 and 15 demonstrate the time-path of optimal policy, m_t, and the corresponding path of infected population and beliefs, assuming that the true infection rate is β_H and β_L. In both

cases, it is assumed that the initial infected population is $I_0 = 0.1$, and $\nu = 0.03$.[41] When $\beta = \beta_H$, the higher π_0, the more accurate the planner's beliefs are. As such, the solid black line indicates the simulation results under "near-perfect" information. As the dashed blue line in Panel (c) illustrates, high structural uncertainty in this case implies an initial under-reaction, followed by a rather prolonged mitigation. This prolonged reaction appears to be a response to a rapid propagation of the infection due to the initial lack of action.

The under-reaction itself seems to be driven mostly by the perceived low cost of inaction, rather than a preference for ignorance. The dash-dotted orange line in Panel (c) supports this: As the initial beliefs—incorrectly—assign more weight to β_L, planner's initial inaction becomes more pronounced, followed by a more severe and more prolonged mitigation compared to both $\pi_0 = 0.9$ and $\pi_0 = 0.5$ cases.

The preceding arguments entirely reverse once the true underlying state is β_L: Here, $\pi_0 = 0.1$ indicates near-perfect information (also illustrated by solid black lines in 15). In this case, an initially misplaced belief on $\beta = \beta_H$ implies a significant overreaction, followed by a relatively rapid re-opening of the economy. An interesting observation is that, significant uncertainty in this case ($\pi_0 = 0.5$) also implies a "rush to reopening," compared to incorrect initial beliefs ($\pi_0 = 0.9$). This result is consistent with Glover et al. (2020)'s criticism of the optimal policy in the U.S.

41. Our results are robust to greater choices for I_0 and/or ν.

FIGURE 14. Time-Path of Policy, and Infected Population and Beliefs under Optimal Policy when $\beta = \beta_H$

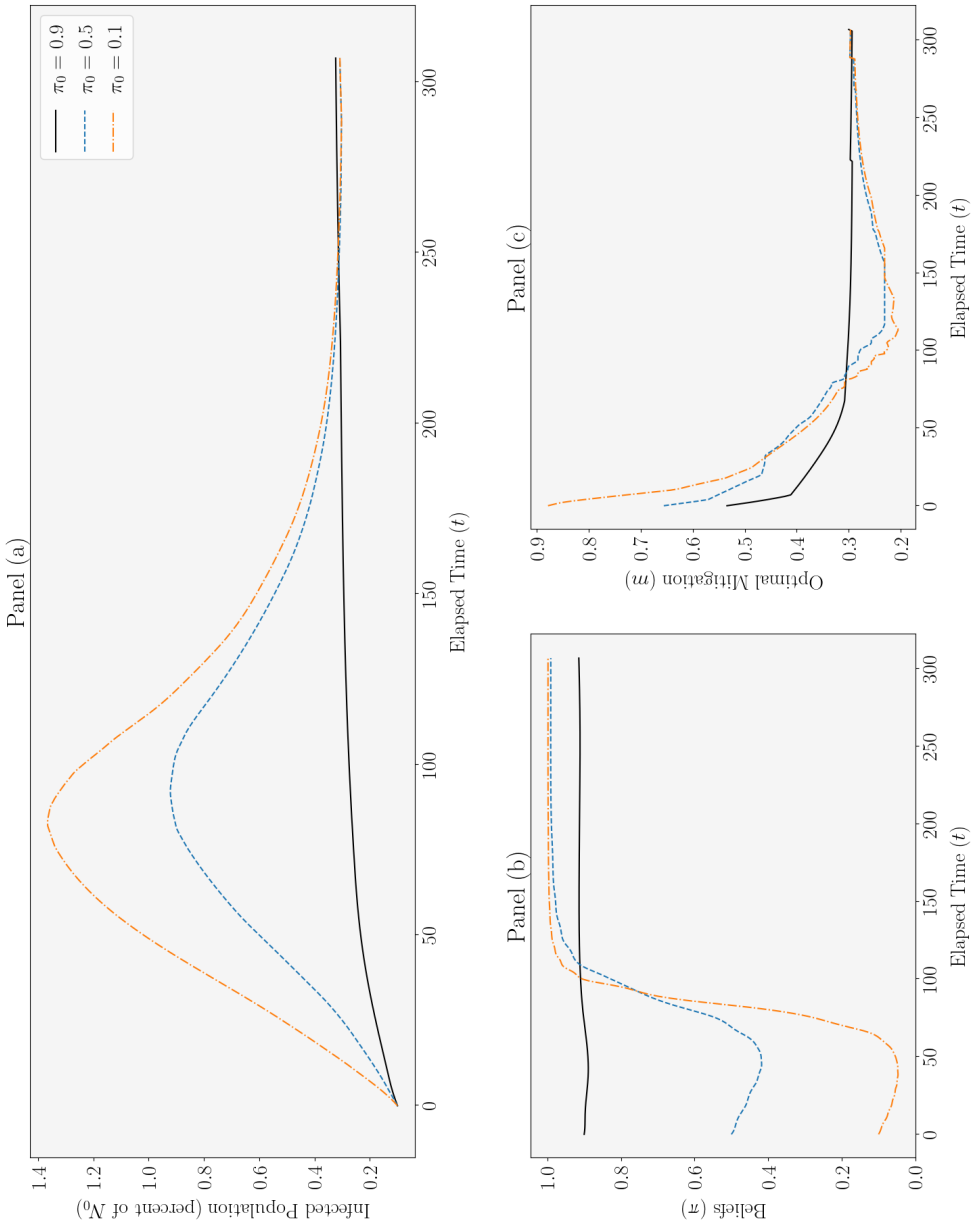

Note. It is assumed $I_0 = 0.1$, $\nu = 0.03$ and $\varepsilon_t = 0$ for all t. Only the first 310 days are drawn. It is assumed that the true infection rate is β_H. Therefore, greater π_0 indicates more accurate beliefs. Discontinuities in optimal paths are artifacts of relatively coarse discrete grids.

FIGURE 15. Time-Path of Policy, and Infected Population and Beliefs under Optimal Policy when $\beta = \beta_L$

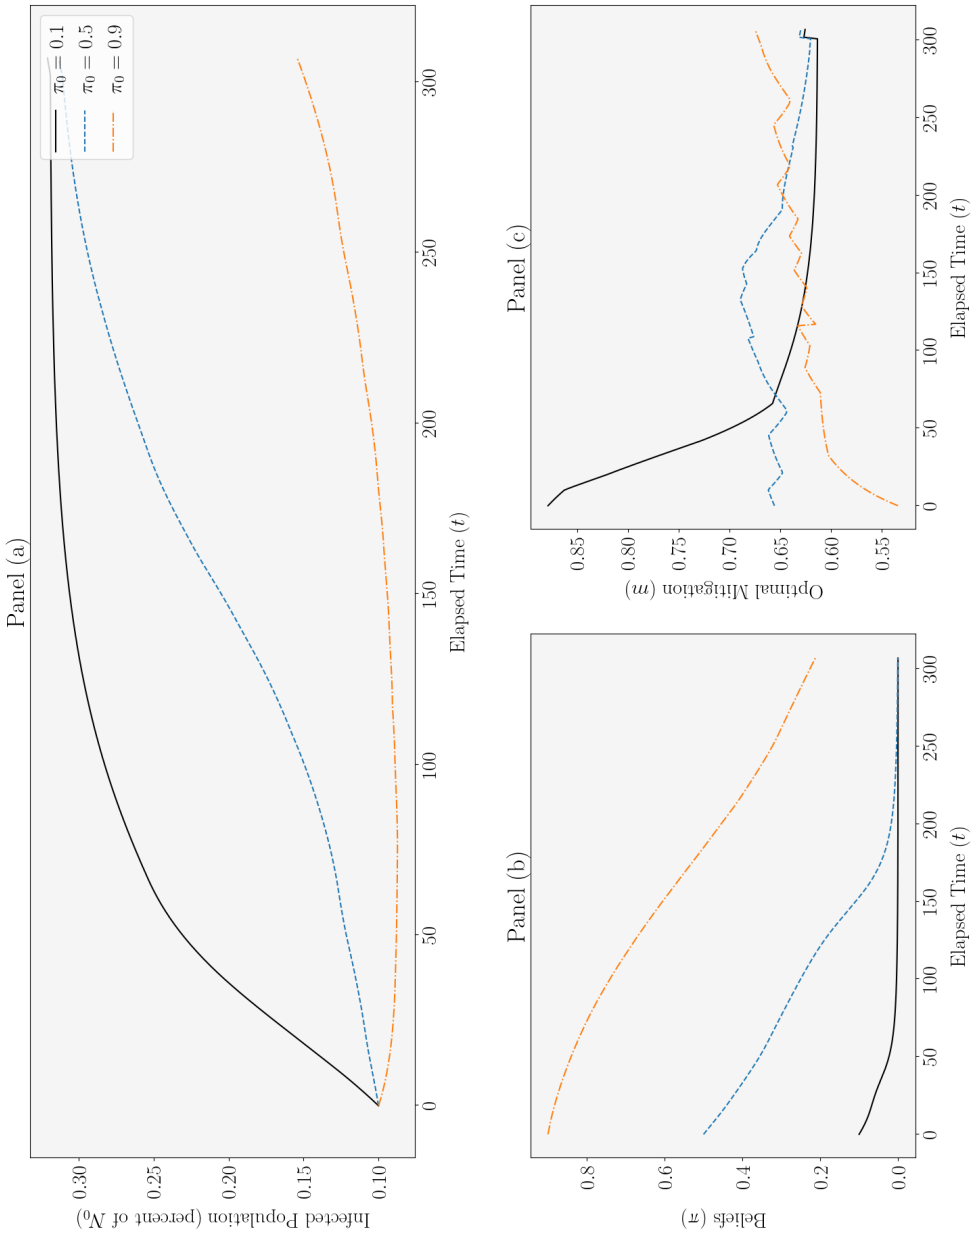

Note. It is assumed $I_0 = 0.1$, $\nu = 0.03$ and $\varepsilon_t = 0$ for all t. Only the first 310 days are drawn. It is assumed that the true infection rate is β_L. Therefore, smaller π_0 indicates more accurate beliefs. Discontinuities in optimal paths are artifacts of relatively coarse discrete grids.