

Optimal Epidemic Control

Martín Gonzales-Eiras, Dirk Niepelt

23-11

October, 2023

DISCUSSION PAPERS

Optimal Epidemic Control*

Martín Gonzalez-Eiras[†]

Dirk Niepelt[‡]

October 4, 2023

Abstract

We develop a flexible single-state model to represent tradeoffs between infections and activity during the early phase of an epidemic. We prove that optimal policy is continuous in the state but discontinuous in the deterministic arrival date of a cure; optimal lockdowns are followed by stimulus policies; and re-infection risk renders laissez faire inefficient even in steady state. Calibrated to the COVID-19 pandemic the model prescribes initial activity reductions of 38 percent. Stimulus policies account for a third of the welfare gains of intervention. Robustness along many dimensions contrasts with sensitivity of the policy prescriptions with respect to the intertemporal elasticity of substitution, activity-infections nexus, and re-infection risk.

JEL codes: D62, I18

Keywords: Epidemic, lockdown, stimulus, logistic model, optimal control, COVID-19

1 Introduction

Sudden epidemiological threats such as the recent COVID-19 pandemic give rise to severe tradeoffs between economic activity and public health. The structure of these tradeoffs is highly uncertain because of varying epidemiological, demographic and economic characteristics, and this complicates an effective policy response. What is needed is a tractable and flexible framework to compare scenarios, identify commonalities, and derive robust policy prescriptions.

*For comments and discussions we thank Alexander Ludwig and seminar participants at ANCE, CEMFI, DTMC COVID-19 Workshop, SED, the University of Copenhagen, UTDT, the VfS Macro meeting, and VMACS.

[†]University of Bologna. Piazza Scaravilli 2, 40126 Bologna, Italy. E-mail: mge@alum.mit.edu. Web: alum.mit.edu/www/mge.

[‡]University of Bern; CEPR. Schanzenekstrasse 1, 3012 Bern, Switzerland. E-mail: dirk.niepelt@unibe.ch. Web: www.niepelt.ch.

No such framework is currently available. The macro-epidemiological literature triggered by the recent pandemic typically employs the classical SIR model (Kermack and McKendrick, 1927) with two epidemiological states. To analyze optimal policy in this environment researchers routinely adopt strong but varying assumptions and resort to numerical simulations.¹ As a consequence, the literature contains a large set of distinct, mostly quantitative findings, commonalities are rare and obscured, and a unified perspective is lacking.

We propose an alternative framework with one epidemiological state variable. The resulting gains in tractability help uncover mechanisms that seem to have gone mostly unnoticed, and the framework’s flexibility as well as the reduced computational burden allow to analyze optimal policy in widely different scenarios. While the one-state simplification comes at a cost—most importantly, orthogonality of policy and long-run health outcomes—we show that this cost is low when analyzing the decisive early stage of an epidemic because the classical SIR model nearly exhibits orthogonality as well.

In fact, when we consider a baseline parameterization and scenario to compare a benchmark SIR-based model with our framework we find essentially the *same* policy prescriptions for the first couple of years. In contrast to the benchmark model, however, our framework can also be employed to analyze a host of other parameterizations and scenarios. Against this background, we view it as a versatile tool for policymakers to flexibly explore the tradeoffs they face at the onset of a coming epidemic.

The generic epidemiological framework we use generalizes the “simple epidemiological model” (Bailey, 1975) and captures the essence of infectious dynamics, namely the interaction between those who have contracted the disease and those who have not but are susceptible. New infections are driven by complementarities between the two groups, and cumulative infections—the single endogenous state variable—approximately follow a logistic law of motion, which accurately approximates infectious dynamics in two-state SIR models.

The economic layer that we superimpose incorporates households and a government. We assume that households derive utility from economic activity, both positive because activity generates consumption and negative because it requires effort. In addition, higher activity increases the risk of getting infected, which is privately costly. Households are fully aware of the aggregate infection dynamics and the risks they face, they end up shouldering the full social costs of infection, and they behave individually rationally. Nevertheless, the privately optimal activity choices fail to fully internalize the social consequences and give rise to “static” and “dynamic” externalities, which reflect wedges similar to those identified in the literature.

It is well known that such wedges call for corrective government action (Gersovitz and Hammer, 2004). More surprisingly, we find that this goes both ways. We characterize the optimal timing not only of lockdowns but also of “inverse lockdowns,” namely interventions that aim at stimulating private sector activity. Intuitively, as the infection spreads,

¹Often, the strong assumptions are needed to minimize numerical problems. For example, Alvarez et al. (2021) assume a counterfactually high share of the initially infected population. Farboodi et al. (2021) solve backwards for the optimal policy by building a chain of piecewise diverging paths and impose that the elasticity of infections with respect to activity equals two.

eventually there are gains from leaving the worst of the epidemic behind. Because individual households do not internalize these social capital gains, the government optimally encourages activity. This can rationalize measures such as monetary easing, temporary sales tax reductions, subsidies, or “return-to-work bonuses” that were proposed or undertaken during the COVID-19 pandemic.

Our second result concerns optimal policy when a cure (a vaccine or effective treatment) is anticipated to arrive deterministically. Such a situation might arise before final regulatory approval of the cure or before donors supply the cure to a country. We prove that the strictness of the optimal lockdown is discontinuous in this case: When the arrival date is near, the optimal policy keeps total infections in check until the cure arrives. But if the date surpasses a critical threshold the optimal lockdown is very light or the government even imposes an inverse lockdown. Intuitively, as the time until the arrival date increases, the costs of curtailing activity by more and more, for longer and longer become prohibitive and it is optimal to “give up” on infections.

Our third result applies when infection and subsequent recovery does not confer immunity. Some households suffer recurrent infections in this case and the economy converges to a policy dependent steady state with strictly positive infection flows. We show that generically, laissez-faire infection dynamics in this endemic-disease scenario are inefficient not only in the transition but also in steady state.

The fourth set of findings comprises quantitative results. We exploit the fact that our framework and extensions thereof can easily be numerically solved and simulated. Crucially from the epidemiological point of view, we solve for the optimal policy under general assumptions about the elasticity of the activity-infection nexus—unlike the literature, which typically needs to assume that this elasticity equals one or two. This allows to account for the mixed evidence on this elasticity ([Hethcote, 1989](#)) and to tailor the model and its policy prescriptions to regional (e.g., rural vs. urban) or other disparities. We also study the role of other key parameter values (e.g., the intertemporal elasticity of substitution) and of the cost function associated with infections, allowing for “flattening-the-curve” motives as well as learning effects. We consider the consequences of virus mutations, enhancements in therapy or test-and-trace technology, or multiple infection waves. We study the effects of restricting the government’s instruments such that inverse lockdowns become infeasible—and find that this is quite costly. And we consider the consequences of information asymmetries and heterogeneity by allowing for publicly or privately observed infection status.

Calibrated to match U.S. COVID-19 infection data the baseline model implies that starting from mid March 2020 economic activity should have immediately been reduced by 38 percent. This would have increased welfare relative to laissez-faire by about three and a half percent of lifetime consumption. Across all the alternative specifications and scenarios we consider these implications are surprisingly robust. There are only three factors that the model prescriptions are very sensitive to: The intertemporal elasticity of substitution, the nonlinearity of the activity-infection nexus, and re-infection risk. Each of these factors is positively associated with optimal strictness.

As noted before, our model replicates the predictions of a benchmark two-state SIR model when we adopt the same functional form assumptions and calibration and rule

out all the non-standard elements our framework allows for. Figure 1 plots the predicted optimal activity level in Farboodi et al. (2021) and in our framework against the share of the infected population (y) during the first two years of the epidemic according to the Farboodi et al. (2021) calibration.² Evidently, the two predictions essentially coincide. That is, our tractable and flexible framework replicates the results in the literature while generating new conceptual insights and allowing for the analysis of a set of much richer scenarios.

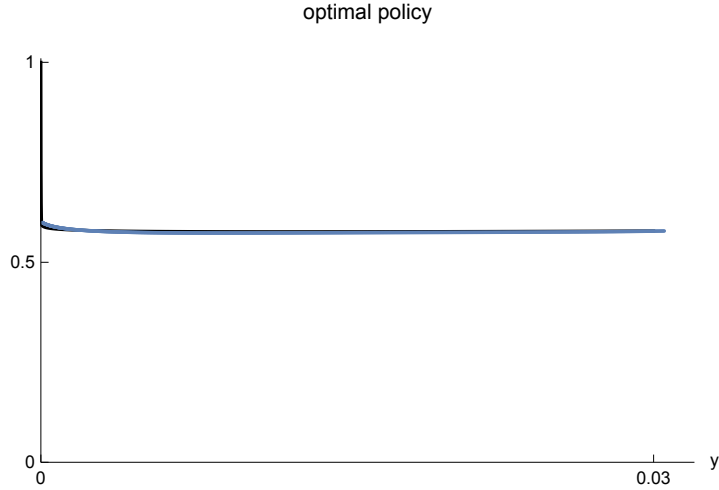


Figure 1: Optimal policy: Logistic model vs. two-state SIR model.

Our final contribution is of a technical nature. The derivation of analytical results as well as numerical solution and simulation in the macro-epidemiological context is challenging since infection dynamics make the Hamiltonian of the control problem non-convex (e.g., Calvia et al., 2023). Although our framework preserves the fundamental nonlinearity of infection dynamics we can show, unlike contributions in the literature, that the value function of the optimizing government is differentiable over the entire state space. This implies that optimal policy is continuous in the state and that standard numerical methods robustly identify the approximate solution of the government’s program.

Related Literature Workhorse epidemiological models due to Kermack and McKendrick (1927) and Bailey (1975) are reviewed, e.g., in Hethcote (1989) and Hethcote (2000). The pandemic has generated broad interest in the intersection of epidemiological dynamics and economic cost-benefit analysis. Early contributions include Atkeson (2020) and Eichenbaum et al. (2021). Alvarez et al. (2021) compute the optimal lockdown policy when there is a rationale to flatten the infection curve in order to relax health care system capacity constraints. Farboodi et al. (2021) argue that in equilibrium and under the optimal policy the effective reproduction number always remains close to unity. Kaplan et al. (2020), Acemoglu et al. (2021), and Ellison (2020) analyze the implications of heterogeneity, including differential costs of reduced activity, asymmetric infection dynamics due

²See appendix A for a detailed discussion.

to “super spreaders,” age-dependent fatality rates, or welfare losses due to nontargeted measures. [Giannitsarou et al. \(2021\)](#) analyze immunity loss and demographic dynamics. Taking spatial aspects into account [Bisin and Moro \(2022\)](#) show how local interactions give rise to matching frictions and local herd immunity effects; [Fajgelbaum et al. \(2021\)](#) characterize optimal policy in a related setting.

Most of this work focuses almost exclusively on numerical analyses, with [Toxvaerd \(2020\)](#), [Gonzalez-Eiras and Niepelt \(2020a\)](#), [Abel and Panageas \(2020\)](#), and [Miclo et al. \(2022\)](#) constituting some notable exceptions.³ Our paper combines novel analytical results with numerical simulations.⁴ [Garibaldi et al. \(2020\)](#) highlight a dynamic externality similar to the one in our model; we show that dynamic externalities necessarily start negative and eventually become positive.

Outline The remainder of the paper is organized as follows. We lay out the model in section 2 and present the conditions characterizing equilibrium and first best in section 3. The theoretical analysis is contained in section 4 and the quantitative analysis in section 5. Section 6 concludes. Auxiliary discussions are relegated to appendices with proofs of the lemmas and propositions in the main text collected in appendix B.

2 The Model

We consider an infinite horizon economy with households and a government. Time is continuous, $t \geq 0$.

2.1 Epidemiology

We use an epidemiological framework that is closely related to several canonical models in the epidemiological literature: The SIR model due to [Kermack and McKendrick \(1927\)](#), a modified SIR and the simple epidemic model (the SI model) due to [Bailey \(1975\)](#), and the SIS model derived from it.⁵ Our framework incorporates one endogenous state variable (rather than two in the typical SIR model), possibly time as a second state variable (unlike SIR and SIS models), as well as economic activity (unlike SIR and SIS models).

2.1.1 Dynamics

The population of mass one consists of mass x “pre-infection” (for short: “pre”) households; mass $y = \bar{y} - x$ “post-infection” (“post”) households; and mass $1 - \bar{y}$ “neutral”

³[Toxvaerd \(2020\)](#) characterizes privately optimal social distancing; [Gonzalez-Eiras and Niepelt \(2020a\)](#) characterize optimal lockdown policies; and [Abel and Panageas \(2020\)](#) characterize the optimal steady state in a SIR model with vital dynamics. [Miclo et al. \(2022\)](#) derive the optimal policy under a capacity constraint.

⁴[Acemoglu et al. \(2021\)](#) and [Farboodi et al. \(2021\)](#) also consider the deterministic case but do not discuss the result that optimal policy may be discontinuous.

⁵The “S,” “I,” and “R” in SIR, SI, and SIS stands for “susceptible,” “infectious,” and “removed,” respectively. See [Hethcote \(1989\)](#) and [Hethcote \(2000\)](#) for an overview over epidemiological models of infectious diseases.

households. Members of the post group have been infected in the past; members of the pre group might be infected in the future; and members of the neutral group cannot become infected, for instance because they are immune.

Initially, at time $t = 0$, the population shares of the pre and post groups are given by $x_0 = \bar{y} - y_0 > 0$ and $y_0 > 0$, respectively. While the infection status of neutral households never changes households in the pre group catch the disease according to a generalized logistic law of motion,

$$\frac{dy}{dt} = f(y, a) \equiv g(a)\beta y\bar{y} \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right), \quad 0 < \beta, \omega < \infty. \quad (1)$$

Accordingly, the share of the pre group changes by $dx/dt = -dy/dt$. Variable a in equation (1) denotes an index of economic activity and function g represents the activity-infection nexus, that is, g is positive and strictly increasing. Parameters β and ω represent epidemiological characteristics.

According to equation (1) the post share, y ; the share \bar{y} ; and the infection rate, $g(a)\beta$, (but not time) determine the speed at which the shares of pre and post households change. With $\omega = 1$ the law of motion is symmetric about $\bar{y}/2$; $\omega \neq 1$ introduces skewness. Starting from $y_0 > 0$ and with $g(a) > 0$ the population share y is strictly increasing over time and converges to \bar{y} ; conversely, the share of the pre group is strictly decreasing and converges to 0. In an extension covered in subsection 4.3 we allow for members of the post group to return to the pre group, for instance because infection does not confer permanent immunity, but in the baseline analysis we abstract from this possibility.

We denote a solution to equation (1) conditional on a control path \mathbf{a} and an initial value y_0 by $y(t; \mathbf{a}, y_0)$; when there is no danger of confusion, we drop the two latter arguments. When activity is constant at level a then⁶

$$y(t) = \frac{\bar{y}}{\left(1 + e^{-\omega g(a)\beta \bar{y} t} \left(\left(\frac{\bar{y}}{y_0}\right)^\omega - 1\right)\right)^{1/\omega}}. \quad (2)$$

With Poisson arrival rate ν a “cure arrives”—the disease and its consequences (described below) disappear and β drops to zero. For example, this might be due to medical progress or the development of a vaccine. The same cure also arrives deterministically, at time T , which may be infinite.

2.1.2 Costs of Infection

Infections are transitions of households from the pre to the post group. They generate social costs, for example because the health care system requires resources, utility is foregone, or infections cause harm. These costs are given by

$$\psi f(y, a), \quad (3)$$

where $\psi > 0$ reflects unit costs. In extensions covered in subsection 5.2 we allow for non-constant ψ in order to analyze learning or congestion effects. In the former case,

⁶See for example [Hethcote \(1989\)](#) for the case of $g(a) = \bar{y} = \omega = 1$.

the unit costs decrease with cumulative case numbers. In the latter case, the unit costs increase with infection flows, for instance due to capacity constraints in the health care sector that lead to a deterioration of care quality and increased fatality rates, generating a motive to “flatten the curve.” In the baseline analysis we let ψ be constant.

2.1.3 Discussion

The law of motion (1) nests a range of well known epidemiological models, augmented by an effect of economic activity on the infection rate, see the discussion in appendix C. Relative to the SIR model our framework eliminates one state variable by combining the currently infected and removed groups in the SIR model into a single group. Importantly, this does not undermine the ability to represent societal costs of infection or death: To represent these costs it suffices to account for the *flow* of newly infected individuals from the pre to the post group and to associate costs with this flow, as we do.⁷

However, the single state variable in our framework does imply that the steady state in the absence of the arrival of a cure, \bar{y} , is exogenous while it is endogenous in the SIR model.⁸ As we demonstrate in appendix A the endogeneity of the steady state in the SIR model is irrelevant for optimal policy choices during the initial phase of the epidemic. For example, we show that during that phase prescriptions of the logistic model and SIR based analyses essentially coincide. We conclude from this that during the initial phase of an epidemic our framework constitutes a useful approximation to the conclusions one would draw (if one could) from a SIR-based analysis.⁹

This holds especially true when we focus on the model predictions in state space rather than the time domain. We have experimented with alternative laws of motion and have found that modifications can have strong effects on the predicted path of y and thus a over time without substantially changing the link between a and y . Against that background we focus on the state-space implications of the model. (Of course, given these state-space results time series predictions follow immediately.)

Formally, we summarize the epidemiological part of the model as follows:

Assumption 1. A cure arrives with Poisson arrival rate $\nu \geq 0$ and deterministically at time T , which may be infinite. Admissible activity levels are $A = [\underline{a}, 1]$ with $0 < \underline{a} < 1$. Before a cure arrives the law of motion $f : [0, \bar{y}] \times A \rightarrow \mathbb{R}^+$ given in equation (1) determines epidemiological dynamics. Function $g : A \rightarrow \mathbb{R}^+$ is strictly increasing, smooth, and weakly convex. Social costs of infection are given by (3).

The lower bound on A can be arbitrarily close to zero. Both bounds on A are not binding but will simplify proofs.

⁷The law of motion (1) does not explicitly account for deaths. The implications for epidemiological dynamics are negligible when death rates are small as we assume. See the discussion in footnote 11.

⁸Because of the random arrival of a cure the long-run share of the population that avoids infection is endogenous to the choice of activity level, both in SIR and our framework.

⁹For our COVID-19 based calibration with a basic reproduction number $\mathcal{R}_0 = 2.4$ a conservative lower bound for this initial phase is a cumulative infection rate of 3.5 percent of the population. Simulation results show that for higher basic reproduction numbers the initial phase extends over a larger part of the state space.

2.2 Economics

2.2.1 Households

A household i chooses activity a_i over time in order to maximize an intertemporal objective which accounts for the immediate economic effects of activity and for the costs of infection that the household privately bears. Households take aggregates as well as the law of motion (1) as given and discount the future at rate ρ .

The immediate economic effects are represented by an indirect utility function, u , that depends on the individual choice, a_i . Utility reaches a maximum at the first-best level in the absence of infection dynamics, a^* . Without loss of generality, we normalize u and g such that $a^* = 1$, $u(a^*) = 0$, and $g(a^*) = 1$.

Assumption 2. Function $u : A \rightarrow \mathbb{R}$ is twice differentiable and strictly concave and satisfies $\infty > u'(\underline{a}) > u'(a^*) = u'(1) = 0$, $u(a^*) = 0$. The discount rate $\rho \geq 0$; when $T = \infty$, then $\rho + \nu > 0$. Function g satisfies $g(a^*) = 1$.

The costs of infection a household privately bears are the product of two terms: The social costs given in (3) and the factor $\xi(a_i, a)$ that reflects the household's activity level and the aggregate level.

Assumption 3. Function $\xi : A^2 \rightarrow \mathbb{R}^+$ is given by $\zeta a_i/a + (1 - \zeta)$, $\zeta \in [0, 1]$.

According to assumption 3 a household bears full social costs in equilibrium. But it internalizes only a fraction ζ of marginal costs and perceives its private marginal costs to depend on own activity relative to aggregate activity; ζ may fall short of unity because of health insurance and non-contractable precautionary behavior or many other reasons. The functional form assumption 3 can substantially be relaxed.¹⁰

Households are representative, rendering y the single endogenous state variable. Representativeness could reflect, for example, that symptoms do not generate much relevant information about their cause or that behavior is unresponsive to the posterior (e.g., due to strong ambiguity aversion) such that symptoms do not trigger changes in behavior. It could also reflect that some households do change behavior but their share constitutes a negligible part of the population.¹¹ Whatever the reasons, independently of infection status all households bear the same share ξ of social costs, which equals unity in equilibrium; and all pre, post, and neutral households (who do not know that they are neutral) choose the same activity level.

The assumption that households behave symmetrically is a plausible approximation in the context of many epidemics, including the COVID-19 pandemic.¹² However, it is clearly

¹⁰See [Gonzalez-Eiras and Niepelt \(2020b\)](#).

¹¹In the law of motion (1) we treat the death rate δ say as negligible. If we explicitly incorporated δ then the population would shrink over time by the measure δy and the payoff from activity would be scaled by $1 - \delta y$. The programs we analyze subsequently would remain unchanged except that β in the law of motion would be scaled by $1 - \delta \approx 1$ and u and U^* by $1 - \delta y \approx 1$.

¹²Many individuals infected with COVID-19 remained asymptomatic or showed only mild symptoms. When only few tests were administered many infected individuals necessarily behaved like individuals who had not been infected.

less plausible in other cases where the symptoms of an infection are more pronounced or easier to differentiate. To account for this possibility we analyze in subsection 5.9 the implications of privately or publicly observable infection status.

2.2.2 Government

Policy makers have instruments at their disposal to control economic activity, for instance by imposing social distancing measures, closing non-essential businesses, mandating other lockdown measures or, in contrast, stimulating activity. Using these instruments the government faces the same program as a social planner.

In subsection 5.5 we analyze the consequences of instrument restrictions, which imply that the government faces a more limited choice set than a social planner.

2.3 Functional Form Assumptions and Calibration

To sharpen analytical results we sometimes impose the preference assumption

$$u(a) = \ln(a) - a + 1,$$

which is consistent with the normalization $a^* = 1$ and implies $u(a^*) = 0$. Our preferred interpretation is that activity yields strictly concave benefits (e.g., from consumption) and linear costs (e.g., from providing effort). In subsection 5.6 we also solve the model under the assumption that the benefit function exhibits stronger curvature.

Moreover, we often assume

$$g(a) = a^n, n \in [1, 2].$$

This specification of the activity-infection nexus allows for both constant and increasing returns to scale and accounts for the mixed epidemiological evidence on the elasticity of infections with respect to activity (Hethcote, 1989).

Assumption 4. Preferences and the activity-infection nexus are given by $u(a) = \ln(a) - a + 1$ and $g(a) = a^n, n \in [1, 2]$, respectively.

Throughout the paper we use numerical simulations to illustrate some of the results.¹³ The simulations make use of the functional form assumptions described above and are based on parameter values that we calibrate to match properties of the COVID-19 pandemic in the United States. Our unit of time is a day and $t = 0$ corresponds to mid March 2020. Accordingly, we set $\rho = -\ln(0.95)/365$ (five percent annual discount rate) and $\nu = 1/(365 * 1.5)$ (one-and-a-half years expected duration until a substantial part of the population is vaccinated or otherwise protected).¹⁴

In appendix D we describe in detail how we calibrate the remaining parameters. Based on information about parameter values in the canonical SIR model we let $y_0 = 0.1148$.

¹³Mathematica code is available upon request.

¹⁴See, e.g., Alvarez et al. (2021). The probability that a cure arrives before time t equals $1 - e^{-\nu t}$; the expected duration thus equals $\int_{t=0}^{\infty} t \nu e^{-\nu t} dt = \nu^{-1}$.

10^{-3} , $\beta = 0.8346 \cdot 10^{-1}$ (corresponding to an infection rate in the SIR model (at regular activity level) of 0.1333), $\omega = 0.6662$, and $\bar{y} = 0.8786$. To calibrate ζ and the social cost parameter ψ we use estimates of expected health care and mortality costs as well as households’ willingness to pay to eliminate COVID-19 induced mortality risk. We assume that households fully internalize mortality risk but not the social costs of health care implying a social cost parameter $\psi = 222.8$ and an internalization rate $\zeta = 0.8266$.¹⁵ Table 1 summarizes the baseline calibration.

Parameter	Value
ρ	$0.1405 \cdot 10^{-3}$
ν	$0.1826 \cdot 10^{-2}$
y_0	$0.1148 \cdot 10^{-3}$
β	$0.8346 \cdot 10^{-1}$
ω	0.6662
\bar{y}	0.8786
ψ	$0.2228 \cdot 10^3$
ζ	0.8266

Table 1: Baseline calibration. See the text and appendix D for explanations.

For the effect of activity on infections, $g(a) = a^n$, we choose the quadratic specification as a baseline, $n = 2$. We do this for two reasons. First, because it is the most widely used assumption in the literature; our choice therefore allows to compare model predictions. Second, because the implications of the model for intermediate values of n , as suggested by epidemiological research (Hethcote, 1989), resemble those of the model with $n = 2$ more closely than those with $n = 1$, as we show below.

3 First Best and Equilibrium

Let $U^* \equiv u(a^*)/\rho$ denote the household’s and the government’s values in the absence of infections, when first-best activity is chosen permanently. This value is attained once the cure arrives or in the limit when all households have gained immunity, $y = \bar{y}$. Prior to attaining U^* , at time $t < T$, the state in the programs of the government or a household is given by (y, d) with $d \equiv T - t$ (for “duration” until T). When $T = \infty$ (no deterministic arrival of a cure) the state only includes y . We proceed using notation for the finite- T case but note how the results translate to the infinite- T case.

¹⁵We rely on parameter estimates by Atkeson (2020), Bartsch et al. (2020), Ferguson et al. (2020), Hall et al. (2020), and Menachemi et al. (2020).

3.1 Government Program

An admissible control path from time 0 to time T is a measurable function $\mathbf{a} : [0, T] \rightarrow A$. Let \mathcal{A} denote the set of such admissible paths and let $V : [0, \bar{y}] \times [0, T] \rightarrow \mathbb{R}$ denote the government's value function prior to the arrival of a cure. At time 0 the value function satisfies

$$V(y_0, T) = P^* + \sup_{\mathbf{a} \in \mathcal{A}} \int_0^T \{u(a) - \psi f(y, a)\} e^{-(\rho+\nu)t} dt \quad \text{s.t. (1),} \quad (4)$$

where y and a arguments are evaluated at $y(t; \mathbf{a}, y_0)$ and $\mathbf{a}(t)$, respectively. When $T = \infty$ equation (4) holds with $V(y_0, T)$ replaced by $V(y_0)$.

The first term on the right-hand side of equation (4), P^* , depends on parameters including T . It represents the present value of U^* realizations either at time T , or after Poisson shocks that occur before time T .¹⁶ The second, integral term represents the probability weighted present value of utility from economic activity net of infection costs before a cure arrives. Note that the upper bound on A that we imposed in assumption 1 is not binding: Not only is marginal utility negative for $a > 1$ but higher activity also speeds up infections, which is costly because of discounting (assumptions 1 and 2).

As a preliminary step to establish differentiability of the value function and continuity of the policy function we show that the Dynamic Programming Principle holds and V is the unique viscosity solution of the Hamilton-Jacobi-Bellman (HJB) equation associated with the optimization problem. Letting $D_y V$ denote the gradient of the value function with respect to y (it is unknown at this point whether the derivative, V_y , exists) and V_d the partial derivative with respect to duration we have:¹⁷

Lemma 1. Under assumptions 1 and 2 the Dynamic Programming Principle holds and V is the unique bounded viscosity solution of the associated HJB equation. When $T < \infty$, the HJB equation reads

$$(\rho + \nu)V(y, d) = \sup_{a \in A} \{u(a) - \psi f(y, a) + f(y, a)D_y V(y, d)\} - V_d(y, d) + \nu U^* \quad \text{s.t. (1)}$$

with initial condition $V(y, 0) = U^*$; moreover, $V(y, d) < U^*$ for all $(y, d) \in (0, \bar{y}) \times (0, T]$ and V is Lipschitz continuous. When $T = \infty$, the same HJB equation holds with $V(y, d)$ replaced by $V(y)$ and boundary conditions $V(0) = V(\bar{y}) = U^*$; moreover, $V(y) < U^*$ for all $y \in (0, \bar{y})$ and V is Hölder continuous with exponent $\min[\frac{\rho+\nu}{\beta \bar{y} \psi \max[\omega, 1]}, 1]$.

¹⁶The probability of no Poisson shock up to time T equals $e^{-\nu T}$ and the probability density of a first shock after duration $t < T$ equals $\nu e^{-\nu t}$. It follows that P^* is given by

$$U^* \cdot \left\{ e^{-\nu T} e^{-\rho T} + \int_0^T \nu e^{-\nu t} e^{-\rho t} dt \right\} = U^* \left\{ e^{-(\rho+\nu)T} + \frac{\nu}{\rho + \nu} (1 - e^{-(\rho+\nu)T}) \right\}.$$

¹⁷A continuous function h on the domain $X \subset \mathbb{R}$ is Lipschitz continuous if there exists a real constant L such that $|h(x_1) - h(x_2)| \leq L|x_1 - x_2|$ for all $x_1, x_2 \in \mathbb{R}$. It is Hölder continuous with exponent $k > 0$ if there exists a constant $C \geq 0$ such that $|h(x_1) - h(x_2)| \leq C|x_1 - x_2|^k$ for all $x_1, x_2 \in \mathbb{R}$. When $k < 1$ Hölder continuity is a strictly weaker requirement.

In appendix E we briefly review the notion of viscosity solutions of nonlinear partial differential equations (Crandall and Lions, 1983). The left-hand side of the HJB equation represents the required return of following the optimal trajectory, and the right-hand side represents the dividend and capital gains components of that return. The dividend component contains the immediate net benefit of activity while the capital gains component reflects the change in value due to the evolution of the state or a sudden arrival of the cure.

Suppose that T is finite and consider a state (y, d) at which V is differentiable such that the derivative $V_y(y, d)$ replaces the gradient $D_y V(y, d)$ (see appendix E). The corresponding optimal control, $a(y, d)$, then solves

$$a(y, d) = \arg \max_{a \in A} \{u(a) + f(y, a)(V_y(y, d) - \psi)\} \quad \text{s.t.} \quad (1)$$

or

$$u'(a(y, d)) = \frac{g'(a(y, d))}{g(a(y, d))} f(y, a(y, d)) (\psi - V_y(y, d)) \quad \text{s.t.} \quad (1). \quad (5)$$

When T is infinite a parallel condition determines the optimal control $a(y)$. In either case the government trades off losses from reduced activity and benefits of slowing down infections and changing the continuation value. Under the baseline functional form assumption 4 condition (5) reduces to the policy function¹⁸

$$a(y, d) = \begin{cases} \frac{1}{1 + \beta \bar{y} y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V_y(y, d))} & \text{if } n = 1 \\ \frac{-1 + \sqrt{1 + 8\beta \bar{y} y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V_y(y, d))}}{4\beta \bar{y} y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V_y(y, d))} & \text{if } n = 2 \end{cases}. \quad (6)$$

Note that both the $V_y(y, d)$ term and $\omega \neq 1$ introduce asymmetry in the policy function.

As stated before conditions (5) and (6) only hold at points where V is differentiable. The following proposition states that this is the case throughout the state space.

Proposition 1. Suppose that assumptions 1 and 2 hold. The policy function is continuous, the value function is differentiable and the two functions satisfy

$$(\rho + \nu)V(y, d) = u(a(y, d)) - u'(a(y, d)) \frac{g(a(y, d))}{g'(a(y, d))} + \nu U^*. \quad (7)$$

Moreover, $\lim_{y \downarrow 0} a(y, d) = a^*$. When $T = \infty$, these results hold for $V(y)$ and $a(y)$.

Continuity of the policy function and differentiability of the value function over the entire state space are generally hard to prove, see Calvia et al. (2023). We establish these properties by exploiting the tight connection between the law of motion and infection costs as well as the fact that the value function is differentiable almost everywhere. According to the proposition the government's value function is completely determined by u , g and the policy function. Under functional form assumption 4 the proposition implies

¹⁸For $n = 2$ the first-order condition has two solutions. We report the unique solution in $(0, 1]$.

$(\rho + \nu)V(y, d) = \ln(a(y, d)) + (1 - a(y, d))(1 - n^{-1})$; for $n = 1$ the government's value function thus equals the scaled logarithm of the policy function.

Focus for now on the time autonomous case ($T = \infty$). Since $V(0) = V(\bar{y}) = U^*$ and $V(y) < U^*$ for all $y \in (0, \bar{y})$, there exists a $y^{\min} \in (0, \bar{y})$ at which V attains its global minimum. This follows directly from the continuity of the value function. In appendix F we characterize y^{\min} and show that V is strictly convex over the domain $[y^{\min}, \bar{y}]$ and $V'(y) < \psi$. To numerically solve for the government's HJB equation we impose condition (6) as well as the terminal condition $V(\bar{y}) = U^*$ and use finite difference methods in Mathematica, see appendix D.3. We use a parallel strategy to characterize the decentralized equilibrium discussed below.

The top left panel of figure 2 illustrates the resulting value function and the top right panel the government's choice of activity.¹⁹ The figure is drawn subject to the baseline calibration introduced earlier and under the assumption that $g(a) = a^2$. (In the quantitative analysis we conduct various robustness checks and consider alternative specifications of the activity-infection nexus $g(a)$.)




Figure 2: Value function, activity level, infections, and dividend in the government's program.

To gain intuition for the shape of V consider the optimal path of infections displayed on the bottom left of figure 2. Infections are hump-shaped (consistent with the predictions of epidemiological frameworks) because the generalized logistic function driving the evolution of y is S-shaped and the endogenous variation of $a(y)$ does not undo this. The costs of infection thus are hump shaped as well and this affects the dividend component in the government's HJB equation, $u(a(y)) - \psi f(y, a(y))$, which is illustrated in the bottom right panel of the figure.²⁰ The capital gains component, $f(y, a(y))V'(y) + \nu U^*$, is the difference between the required return, $(\rho + \nu)V(y)$, and the dividend component.

¹⁹It is understood that these relationships apply before a cure arrives. Once a cure has arrived a optimally equals unity even if $y < \bar{y}$.

²⁰Dividends are negative because $u(a^*)$ is normalized to zero.

Over most of the state space the slope of V is relatively constant. Since infections are hump shaped this translates into hump shaped capital gains. When infection rates are high, costs of infection and losses from low activity are high as well and the government's value, which represents discounted future costs and losses, rises swiftly. That is, dividends are low and capital gains large when infection rates are high. For small values of y infection rates are tiny and as a consequence, dividends are not much depressed. Capital gains are negative in this part of the state space because the period with low dividends, due to high infection rates and low activity, is approaching.

Most of the literature assumes $g(a) = a^n$ with $n = 1$ or $n = 2$ because the numerical solution strategy described above or similar standard strategies run into problems when $n \in (1, 2)$. In contrast, we can easily solve the model for such intermediate values of n : Exploiting proposition 1 in combination with the optimality condition (6), we can directly solve a differential equation in $a(y)$ with boundary condition $a(\bar{y}) = 1$.

3.2 Decentralized Equilibrium

Unlike the government an individual household takes the law of motion of the state as given. As a consequence it solves a static problem, trading off net economic activity benefits and infection costs. The household's choice conditional on the aggregate state solves

$$a_i(y, d) = \arg \max_{\alpha \in A} u(\alpha) - \xi(\alpha, a(y, d)) \psi f(y, a(y, d))$$

subject to (1) where $a(y, d)$ denotes the aggregate activity choice which the household takes as given. Note that the effect of individual and aggregate activity choices on infections represented by the rightmost term resembles typical specifications in the literature.²¹ In equilibrium individual and aggregate choices coincide and thus satisfy

$$u'(a(y, d)) = \frac{\zeta}{a(y, d)} \psi g(a(y, d)) \beta y \bar{y} \left(1 - \left(\frac{y}{\bar{y}} \right)^\omega \right). \quad (8)$$

Any fixed point $a(y, d) \in A$ that solves this equation constitutes an equilibrium activity choice conditional on the state. This choice has several noteworthy properties. First, unlike the government's activity choice equilibrium activity is symmetric around $\bar{y}/2$ when $\omega = 1$. Second, conditional on y equilibrium activity is independent of time even when $T < \infty$. Third, as is straightforward to show, in our benchmark specification (under assumptions 3 and 4) equation (8) has a unique solution in \mathbb{R}^+ and this solution lies in the interval $(0, 1)$.²² For $n = 1, 2$ this solution is given by

$$a(y, d) = \begin{cases} \frac{1}{1 + \beta \bar{y} y \left(1 - \left(\frac{y}{\bar{y}} \right)^\omega \right) \zeta \psi} & \text{if } n = 1 \\ \frac{-1 + \sqrt{1 + 4 \beta \bar{y} y \left(1 - \left(\frac{y}{\bar{y}} \right)^\omega \right) \zeta \psi}}{2 \beta \bar{y} y \left(1 - \left(\frac{y}{\bar{y}} \right)^\omega \right) \zeta \psi} & \text{if } n = 2 \end{cases}. \quad (9)$$

²¹Under assumptions 3 and 4 and when $\zeta = 1$ the individual perceives health costs to be proportional to $\alpha \cdot a(y, d)$.

²²By letting \underline{a} be sufficiently small we can thus guarantee existence of equilibrium.

An equilibrium activity function consistent with (8) and the law of motion (1) jointly induce equilibrium infection dynamics. They determine the equilibrium value function prior to the arrival of a cure, which we denote by $U(y, d)$. Smoothness of u and g imply that activity is a smooth function of the state as well and this implies that U is differentiable.

4 Analytical Results

In this section we derive several theoretical results and we show that these results apply much more widely than in the specific model analyzed here. Specifically, we characterize static and dynamic externalities and show that the optimal policy does not only involve lockdowns but also inverse lockdowns; we prove that optimal policy is discontinuous when the cure arrives deterministically; and we establish generic inefficiency of steady-state equilibrium activity in an extension with re-infection risk.

4.1 Externalities, Lockdowns, and Inverse Lockdowns

That individual households do not perceive their choices to affect the aggregate state, and thus the continuation value, gives rise to a “dynamic” externality. In addition a “static” externality arises because households do not fully internalize the contemporaneous costs of infection. This is reflected in the first-order conditions (5) and (8), which reveal two differences between the marginal costs of activity as perceived by an individual household and the government. Subtracting the right-hand side of (5) from the right-hand side of (8) and evaluating terms at a common activity level yields $\mathcal{S}(a, y) + \mathcal{D}(a, y, t)$ with

$$\begin{aligned}\mathcal{S}(a, y) &\equiv \psi f(y, a) \left(\frac{\zeta}{a} - \frac{g'(a)}{g(a)} \right), \\ \mathcal{D}(a, y, t) &\equiv f(y, a) \frac{g'(a)}{g(a)} V_y(y, t).\end{aligned}$$

Here, $\mathcal{S}(a, y)$ and $\mathcal{D}(a, y, t)$ represent the static and dynamic externalities, respectively.

Both externalities are proportional to infections, $f(y, a)$, because these drive the costs of infection and change the state variable, which in turn affects the continuation value. The static externality arises for two conceptually distinct reasons. First, because households view the costs of infection they bear to only partially depend on their own activity choice if $\zeta < 1$. Second, because they perceive their private health cost to be linear in their activity choice while at the aggregate level, depending on $g(a)$, the cost function might be strictly convex. In particular, when $g(a) = a^n$, the static externality is proportional to $(\zeta/n - 1)/a$ and is strictly negative when $\zeta < 1$ or $n > 1$.

Figure 3 illustrates the consequences of the externalities. As in figure 2, we focus on the time autonomous case and let $n = 2$; as discussed previously, we calibrate $\zeta = 0.8266$. The solid lines in the figure represent the outcomes implemented by the government and correspond to the schedules in figure 2; the dashed lines represent the equilibrium outcomes.




Figure 3: Value function and activity level in the government's program (solid) and in equilibrium (dashed).

Clearly, the presence of externalities implies that the equilibrium value falls short of the value under the optimal policy. More interestingly, the activity levels displayed in the right panel differ markedly: Initially, the government chooses a lower activity level than in equilibrium, eventually the opposite holds true. The driving force behind this reversal is the capital gains component which only the government internalizes and which is strictly positive for $y > y^{\min}$. When this component is sufficiently strong to compensate for the negative static externality (due to $\zeta < n$) then the equilibrium activity level falls short of the level chosen by the government.

We refer to “lockdowns” or “inverse lockdowns,” respectively, as situations in which the government wishes to depress or stimulate economic activity. Lockdown measures may include for instance stay-at-home-orders, social distancing rules, business closures, or school closures while inverse lockdowns may take the form of stimulation measures such as monetary easing, temporary sales tax reductions, subsidies, or a “return-to-work bonus.” From the preceding analysis we know that the static and dynamic externalities require a lockdown when $\psi(\zeta/n - 1) + V_y(y, t) < 0$ and an inverse lockdown when the reverse inequality holds (assuming $g(a) = a^n$).²³

Let y^c denote the smallest value of y (if it exists) at which the total externality equals zero, $\psi(\zeta/n - 1) + V_y(y^c, t) = 0$. The following proposition establishes that the optimal policy involves an immediate lockdown period followed by at most one inverse lockdown period:

Proposition 2. Under assumptions 1 and 3 and if $g(a) = a^n$ and $\zeta \leq n$, lockdowns occur as follows:

- i. Starting from small y , the government immediately imposes a lockdown;
- ii. under assumption 2 and if $T = \infty$, $\frac{g(a^*)\beta\bar{y}\omega}{\rho + \nu + g(a^*)\beta\bar{y}\omega} > 1 - \zeta/n$, the government relaxes the lockdown at the unique y^c and immediately imposes an inverse lockdown.

²³Note that the parameter ν does not directly enter the condition for an (inverse) lockdown. It matters indirectly, however, because it shapes the value function and thus its derivative.

The first part of proposition 2 is consistent with the fact that during the COVID-19 pandemic many governments imposed lockdown measures early on. The second part appears more surprising but it is consistent with various observed stimulus measures such as temporary tax reductions, employment or consumption subsidies. In subsection 5.5 we analyze the consequences of constraints that make it impossible for the government to impose an inverse lockdown.

Note that the basic intuition underlying the reversal result is very general: Since an epidemic generates costs the value during the transition is lower than when the transition is completed; that is, at some point, society experiences capital gains. These capital gains arise due to the change of an aggregate state variable (or many such state variables) which a household takes as given; that is, the capital gains are not internalized by households. As long as the capital gains are sufficiently large to outweigh negative static externalities the reversal result thus follows.

4.2 Deterministic Arrival of a Cure: Policy Discontinuity

In some contexts a deterministic arrival date of a vaccine or other cure appears more plausible than a stochastic one. For example, when several promising candidates for a vaccine undergo final clinical trials, as was the case in the context of COVID-19 in fall of 2020, it may be near certain that at least one candidate will soon be approved by regulatory bodies and available for distribution. Or, as was similarly the case at the time, a developing country that is dependent on international assistance might anticipate that it will have to wait for a certain period until a vaccine will be supplied by donor institutions and distributed domestically.

To represent an environment with a deterministic arrival date, we let $T < \infty$ (making time a second state variable) and $\nu = 0$. Since we are only interested in policy and not the value function it suffices to use Pontryagin's maximum principle and formulate the government's program using the Hamiltonian²⁴

$$\mathcal{H}(a(t), y(t), t) = u(a(t)) + (\mu(t) - \psi)g(a(t))\beta\bar{y}y(t) \left(1 - \left(\frac{y(t)}{\bar{y}}\right)^\omega\right),$$

where $(a(t), y(t))$ denote the control and state at time t and $\mu(t)$ denotes the (present value) multiplier associated with the law of motion. An optimal plan satisfies the boundary condition $\mu(T) = 0$ as well as the conditions $\mathcal{H}_a(a(t), y(t), t) = 0$, $\mathcal{H}_y(a(t), y(t), t) = -\dot{\mu}(t) + \rho\mu(t)$, and the law of motion. In contrast, the equilibrium activity path solves

$$u'(a(t)) = \frac{\zeta}{a(t)}\psi g(a(t))\beta\bar{y}y(t) \left(1 - \left(\frac{y(t)}{\bar{y}}\right)^\omega\right)$$

and the law of motion.

²⁴From the analysis in section 3.2 the decentralized choice of activity is independent of ν or duration (and thus of T). The multiplier $\mu(t)$ in the Hamiltonian corresponds to the partial derivative $V_y(y(t), T-t)$ such that $\dot{\mu}(t) = -V_{yd}(y(t), T-t) + V_{yy}(y(t), T-t)g(a(t))\beta\bar{y}(1 - (1+\omega)(y(t)/\bar{y})^\omega)$. From proposition 1 the value function is differentiable.

The following proposition establishes that the government’s optimal policy may not be continuous in T . We prove this for the special case in which $\rho = 0$ because it allows for a sharper characterization:

Proposition 3. Under assumption 1 and if $\rho = \nu = 0$ and $T < \infty$, the government’s optimal activity choice is constant until $t = T$. If, moreover, $u(a) = \ln(a) - a + 1$, $g(a) = a^n$, $n = 1, 2$, and $y(0)$ is sufficiently small (as defined in the proof) the optimal activity is discontinuous in T .

The first part of proposition 3 implies that the government’s program reduces to the choice of a fixed activity level. With some abuse of notation we denote this constant level by $a(T)$ to emphasize that it depends on the arrival date. The second part states that $a(T)$ is not continuous in the arrival date. Starting from a low value of T , the optimal activity choice falls as T increases; at the same time, the stock of post households at the time when the cure arrives, $y(T)$, increases with T . As T increases further, the costs of curtailing activity eventually become so large that it is no longer optimal to impose a severe lockdown in order to keep $y(T)$ in check. Instead it becomes optimal to choose a much higher activity level and to “give up,” i.e., to accept higher infection numbers.²⁵

To understand this result note that cumulative infections evolve S-shaped as a function of T and reductions in activity suppress new infections and thereby delay the “peak infection period” during which y increases most rapidly; that is, increases in activity shift the S-shaped $y(T)$ -function to the left (and rescale it). For fixed T the marginal cost of activity due to higher cumulative infections therefore is non-monotone: At low levels, an increase in activity has minor costs; for intermediate levels, the marginal cost is larger; and for high levels the marginal cost is minor again. This non-monotone marginal cost of activity contrasts with monotonically decreasing (and convex) marginal benefits of economic activity.




Figure 4: Constant activity level $a(T)$ (solid) and terminal share of post-group $y(T)$ (dashed) as functions of T .

²⁵The government’s program does not satisfy Arrow’s second-order conditions because the maximized Hamiltonian fails to be concave in y . Accordingly, the first-order conditions are not sufficient for a global maximum.

Figure 4 illustrates the policy discontinuity when we impose the baseline calibration (except for $\rho = \nu = 0$) and $g(a) = a^2$. As T increases, the optimal activity level initially falls before eventually (when T approaches 600) jumping to a value close to unity. Correspondingly, $y(T)$ initially increases before eventually jumping to a value close to \bar{y} . Because S-shaped infection dynamics are a general feature of epidemiological models we expect the essence of proposition 3 to apply much more widely, and indeed find this to be the case in simulations.²⁶

4.3 Lack of Immunity: Generic Steady-State Inefficiency

In the baseline specification households that undergo infection and recovery are protected from re-infection such that herd immunity is eventually attained. But immunity may wane and individuals may repeatedly contract the disease. To study the implications of such lack of immunity we now assume that previously infected persons may join the pre group again and subsequently re-contract the disease. Formally, we replace equation (1) by the modified law of motion

$$f(y, a) \equiv g(a) \beta y \bar{y} \left(1 - \left(\frac{y}{\bar{y}} \right)^\omega \right) - \gamma y, \quad \omega > 0, \gamma \geq 0, \quad (10)$$

which only differs from (1) because of the $-\gamma y$ term. According to equation (10) gross flows into the post pool are determined exactly as in the main model but net flows are smaller because a fraction γ of households in the post group loses immunity and moves back to the pre pool. Unlike in the main model, there is thus no guarantee that y monotonically increases (unless a is constant over time). For $g(a) > 0$, y does not converge to \bar{y} but instead to zero or

$$\bar{y} \left(1 - \frac{\gamma}{g(a)\beta\bar{y}} \right)^{\frac{1}{\omega}},$$

depending on whether $g(a)\beta\bar{y}$ is smaller or larger than γ . We consider the latter case.

In this framework the relationship between γ and the steady-state infection flow $\gamma y(\infty)$ is inverse-U-shaped. Since we view a positive effect of γ on the costs of infections as more plausible (in line with extended SIR models with loss of immunity²⁷) we restrict attention to small values of γ which guarantee that $\gamma y(\infty)$ increases in γ . Specifically, we require

$$\frac{\partial \gamma \bar{y} \left(1 - \frac{\gamma}{g(a(\infty))\beta\bar{y}} \right)^{\frac{1}{\omega}}}{\partial \gamma} > 0 \quad \Leftrightarrow \quad g(a(\infty))\beta\bar{y} > \gamma \frac{1 + \omega}{\omega},$$

where we use the condition $g(a)\beta\bar{y} > \gamma$ and $a(\infty)$ denotes steady-state activity. The following assumption modifies assumption 1 to take these changes into account:

²⁶In a SIR model calibrated to match features of the COVID-19 epidemic (see appendix D, we also impose $\nu = \rho = 0$) we find that optimal activity jumps up at $T \approx 2030$ days, from 0.655 to 0.895, with cumulative infections at T changing from 3.1 to 77.7 percent.

²⁷In an extended SIR model with loss of immunity (a SIR-S model) where $x(t)$, $\iota(t)$, and $z(t)$ denote the stocks of the susceptible, infected, and recovered population, respectively, the steady state satisfies $bx(\infty)\iota(\infty) = c\iota(\infty) = \gamma z(\infty)$, implying $d\iota(\infty)/d\gamma > 0$ (see appendix D).

Assumption 5. Epidemiological dynamics are represented by the law of motion (10). Function g is smooth, positive and strictly increasing, and weakly convex. The social costs of infection are given by (3). Moreover, $g(a(\infty))\beta\bar{y} > \gamma(1 + \omega)/\omega$.

For $g(a) = a^n$ the law of motion (10) implies that in steady state

$$y(\infty) = \bar{y} \left(1 - \frac{\gamma}{a^n(\infty)\beta\bar{y}} \right)^{1/\omega}. \quad (11)$$

Moreover, in the government's program the steady state also satisfies optimality condition (6) evaluated at $(a(\infty), y(\infty))$ as well as the envelope condition subject to $f(y(\infty), a(\infty)) = 0$, while in equilibrium it also satisfies condition (9) evaluated at $(a(\infty), y(\infty))$. Since $a(\infty) > 0$ assumption 5 necessarily is satisfied for sufficiently small γ . We then have the following result:

Proposition 4. Under assumptions 2–5 and $T = \infty$ the steady-state equilibrium activity level is generically inefficient. If, in addition, $u(a) = \ln(a) - a + 1$, $g(a) = a^n$, and $n = 1, 2$, the government's choice of $a(\infty)$ exceeds the equilibrium value if and only if $(\rho + \nu)(n - \zeta) < \zeta\beta\bar{y}\omega$.

According to proposition 4 the government's steady-state activity level necessarily exceeds the equilibrium level if there is no static externality ($\zeta = n = 1$). Even with static externalities ($\zeta < n$) equilibrium activity may be suboptimally low; this is more likely to be the case when there is more concern about the future (low ρ and/or ν), the infection spreads quickly (high β), or a small share of the population is neutral (high \bar{y}). However, in the extreme case where households do not perceive any direct health costs ($\zeta = 0$) equilibrium steady-state activity necessarily is too high. We analyze off-steady-state dynamics in section 5.

5 Quantitative Results

Exploiting the framework's tractability we next turn to its quantitative implications. We allow for endogenous costs of infection, vary the activity-infection nexus and the consumption smoothing motive, introduce stochastic regime change and constraints on the government's use of policy instruments, and analyze re-infection risk as well as publicly or privately observed infection status.

5.1 Baseline Model

We focus on three key statistics to describe the policy implications. First, the share of the infected population at the end of the lockdown, y^c ; in the baseline model this is given by 0.1027. Second, the intensity of the lockdown; against the background of our discussion in subsection 2.1 and appendix A, we report the average activity level until $y(t)$ reaches 3.5 percent or y^c , whatever is smaller. In the baseline model this average, \hat{a} say, equals 0.6238. And third, the welfare gain due to optimal lockdowns/inverse lockdowns. Evaluated at

Scenario	y^c	\hat{a}	$\phi^u - \phi^v$
Baseline model	0.1027	0.6238	0.0348
Endogenous costs: Congestion	0.1383	0.6228	0.0499
Endogenous costs: Learning	0.0787	0.6618	0.0277
Linear activity-infections nexus	0.0322	0.7340	0.0558
Intermediate activity-infection nexus	0.0601	0.6629	0.0384
Regime change: Fall in β	0.1532	0.5633	0.0377
Regime change: Multiple waves	0.3275	0.5509	0.0894
Constraints on policy instruments	0.1188	0.6057	0.0222
Stronger curvature of u	0.0766	0.7853	0.0166
Higher fatality rate	0.1301	0.5310	0.0565
Slower arrival rate of a cure	0.0259	0.7995	0.0191
Re-infection risk	0.4517	0.4411	0.1141
Observability: Public information	0.2653	0.6771	0.0224
Observability: Private information	0.1684	0.6238	0.0460

Table 2: Key statistics for different scenarios. y^c denotes the value of the state at which the lockdown ends; \hat{a} denotes the average activity level during the lockdown or until $y(t) = 0.035$; and $\phi^u - \phi^v$ denotes the welfare gain of the government intervention. Numbers are rounded to four digits. See explanations in the text.

$y(0)$, corresponding to the share of the U.S. population infected in mid March 2020, we find for the baseline model $U(y(0)) \approx -150.3$ and $V(y(0)) \approx -127.1$. The equivalent reduction in consumption, either in equilibrium or under the optimal policy, ϕ^u and ϕ^v respectively, thus solves

$$\begin{aligned} \frac{1}{\rho + \nu}(\ln(a^*(1 - \phi^u)) - a^* + 1) &= -150.3, \\ \frac{1}{\rho + \nu}(\ln(a^*(1 - \phi^v)) - a^* + 1) &= -127.1. \end{aligned}$$

This yields $\phi^u - \phi^v = 0.0348$, i.e., the lifetime consumption equivalent of optimal government equals three and a half percent. We summarize these findings as well as the key statistics of all other scenarios in table 2.

5.2 Endogenous Costs: Congestion and Learning Effects

We relax the assumption that the unit costs of infection are constant and consider two alternative specifications. In the first, we assume that the costs of infection are quadratic in the infection flow, capturing congestion effects that can arise from capacity constraints in the healthcare sector. This change requires recalibration. To simplify comparisons we assume that the change of specification does not alter the total health costs over

the duration of the pandemic (see appendix D.2). The convexity assumption therefore only changes how these costs occur over time. (In subsection 5.7 we analyze how higher total costs affect optimal policy.) In the second specification, we assume that the cost of infection per unit infection flow decreases as the share of the population previously infected increases, reflecting learning effects in healthcare, administration and logistics.

We find that congestion effects imply a longer optimal lockdown duration and larger welfare gains from optimal policy compared to the baseline. In contrast, learning effects shorten the optimal lockdown duration and reduce the welfare gains from policy intervention.

5.3 Modified Activity-Infections Nexus

Next, we modify the activity-infections nexus encoded in function g in equation (1). We consider two alternatives to the baseline specification with $g(a) = a^n, n = 2$. First, the linear case with $n = 1$ in which g exhibits constant returns to scale. And second, an intermediate case with modestly increasing returns, $n = 1.5$. In either case, $n < 2$ dampens the impact of activity reductions on infection flows when activity exceeds one half, and strengthens it when activity is smaller than one half.²⁸ In our analysis the case with $a(y) > 0.5$ is the relevant one. Accordingly, the government perceives a weaker benefit of activity reductions when $n < 2$ and lowers activity by less than in the baseline scenario. In equilibrium, in contrast, households perceive higher marginal costs of activity (proportional to $\zeta g(a(y))/a(y)$) when $n < 2$. As a consequence, households behave more cautiously than in the baseline scenario.

Figure 5 shows the optimal activity levels for $n = 1, 1.5, 2$. As anticipated in subsection 2.3, the optimal policy when $n = 1.5$ is closer to the optimal policy when $n = 2$. When $n = 1$ the optimal lockdown is substantially shorter and less stringent than in the baseline. Due to households' more cautious and the government's less cautious activity choices, the optimal government intervention—in particular the inverse lockdown—has a larger positive welfare effect.

5.4 Stochastic Regime Change: Fall in β and Multiple Waves

Next, we ask how expected changes in the epidemiological environment affect optimal policy. We consider two scenarios. In the first we allow for a permanent reduction in β by fifty percent (reflecting, e.g., improved test, trace, and quarantine strategies) that materializes stochastically, with arrival rate $\mu = 1/90$.²⁹ The HJB equations that apply before the regime change therefore feature additional capital gains terms. In the second scenario we allow for recurrent waves of global loss of immunity (e.g., due to new virus

²⁸Note that $g'(a)$ increases in n when $a > 0.5$ and decreases in n when $a < 0.5$. See equations (6) and (9) for the expressions for a in equilibrium and under the optimal policy when $n = 1$.

²⁹This reflects an expected duration of three months until the regime changes. The U.K. implemented a test-and-trace regime roughly three months into the COVID-19 epidemic. Fetzner and Graeber (2020) conclude that the new regime lowered infection rates. Piguillem and Shi (2022) stress the advantages of strict testing over lockdowns.




Figure 5: Activity level in the government’s program: Specifications with $n = 1$ (dashed, top), $n = 1.5$ (solid), and $n = 2$ (dashed, bottom).

strains) that arrive once per year on average and cause the number of persons in the post group to revert back to $y(0)$. This also introduces new capital gains terms in the HJB equations.³⁰

The upside risk of lower β in the first scenario increases the pre-regime-change value function compared with the baseline and leads the government to wait longer before relaxing the lockdown.³¹ The prospect of multiple waves in the second scenario induces the government to behave even more cautiously. The optimal lockdown lengthens substantially and the welfare gains of optimal government intervention increase.

5.5 Constraints on Policy Instruments

To clarify the importance of lockdown versus inverse lockdown measures we assume next that the government can curtail economic activity but is unable to impose inverse lockdowns. For high values of y the government’s value function then coincides with the *equilibrium* value function, and across the state space the government’s value falls short of the planner’s. Solving for V requires a modified boundary condition based on value matching and smooth pasting conditions.³²

The restriction on inverse lockdown measures implies a longer and slightly stricter lockdown.³³ Although households act too cautiously during the later stage of the epidemic and the government cannot correct this, it thus acts more—not less—cautiously also during the early stage. The welfare gains of the optimal intervention are substantially reduced and this is even more the case when $n = 1$, consistent with the discussion in subsection 5.3. We thus arrive at the surprising result that inverse lockdowns are responsible for a large share of the welfare gains of optimal policy, namely one third when $n = 2$ and almost 90 percent when $n = 1$.

³⁰We could enrich the analysis by letting infection and/or fatality rates vary across waves.

³¹The opposite holds true when the government is concerned about downside risk.

³²At \hat{y} say the government ends the lockdown and would start an inverse lockdown if it were able. We first solve for U and then find V and \hat{y} from the conditions $V(\hat{y}) = U(\hat{y})$, $V'(\hat{y}) = U'(\hat{y})$.

³³In all other scenarios the lockdown ends when the total externality equals zero. This is not the case here.

5.6 Stronger Curvature of u

Next we consider a modified specification of preferences, letting net utility from activity equal $u(a) = -a^{-1} - a + 2$ rather than $u(a) = \ln(a) - a + 1$. That is, we assume that the intertemporal elasticity of substitution for consumption equals one half rather than one as posited in the baseline scenario.

The reduced intertemporal substitutability renders it costlier to lower consumption and activity therefore falls by less than in the baseline scenario, both in equilibrium and under the optimal policy. The optimal lockdown is shorter and the welfare gains due to the optimal government intervention are smaller.

5.7 Higher Fatality Rate and Lower Arrival Rate of a Cure

We also consider the effects of an increase in the fatality rate by fifty percent as well as a lower arrival rate of a cure (increasing the expected duration until discovery of a cure from 1.5 to 5 years), possibly due to lower R&D expenditures for measures to counteract a regional epidemic.

In the first case the government acts more cautiously while the opposite happens in the second case. In contrast, households reduce equilibrium activity only in response to a higher fatality rate; the arrival rate has no effect on their behavior because it does not affect the static tradeoff between activity and infection risk. Relative to the baseline, the benefits of government intervention are higher in the first scenario and lower in the second one.

5.8 Re-Infection Risk

With re-infection risk $\bar{\gamma}$ no longer constitutes a rest point and we thus impose modified boundary conditions to solve for the value functions, see subsection 4.3.³⁴ We assume that immunity is lost on average after a year (Giannitsarou et al., 2021), $\gamma = 1/365$, comparable with the multiple waves scenario discussed earlier. Maintaining all other parameter values we find that the steady state $(a(\infty), y(\infty))$ optimally equals $(0.8754, 0.8151)$ while in equilibrium it is given by $(0.6198, 0.7535)$.

The presence of re-infection risk makes the optimal lockdown substantially longer and stricter and it strongly increases the welfare benefit of optimal government intervention.

5.9 Observable Infection Status

In the baseline model households in the pre, post, and neutral groups behave symmetrically because they are unaware of their infection status. As we discussed in section 2 this is a plausible assumption in the context of some epidemics, including the COVID-19 pandemic, but less so in others. We therefore analyze the consequences of observable infection status including heterogeneous information sets, activity choices, and continuation

³⁴The value function evaluated at steady state equals the capitalized steady-state utility flow from activity net of the steady-state flow costs of infection.

values. Specifically, we assume that a share σ of infected households develop traceable symptoms such that they become aware of their infection status. A share σy of the population thus knows that it is immune.³⁵ It chooses activity a^* and its value equals U^* .³⁶ We consider both public and private information scenarios. Appendix G shows how observability affects the law of motion for infections and contains the analysis of the government's program and the equilibrium.

When the infection status is public information all unaware households in the pre, post, and neutral groups avoid contact with the aware households in the post group. Since households in the pre group only interact with unaware members of the post group the law of motion reads

$$f(y, a) = g(a)\beta\bar{y}(1 - \sigma)y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right). \quad (12)$$

All unaware households have the same activity level a . This law of motion also applies under the Ramsey policy because it is optimal for the government to separate aware households from the rest of the population. In appendix G we analyze the government's program and the equilibrium.

When the infection status is private information³⁷ all post households interact with pre households. Pre households and unaware post households choose the activity level $a(y)$, aware post households choose a^* . The law of motion thus reads

$$f(y, \bar{a}) = g(\bar{a})\beta\bar{y}y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right), \quad (13)$$

where $\bar{a} \equiv (\sigma ya^* + (\bar{y} - \sigma y)a)/\bar{y}$ denotes average activity. Appendix G contains the analysis of the government's program and the equilibrium.

For $\sigma = 0.1$, optimal and equilibrium activity in both the public and private information scenarios are nearly indistinguishable from their counterparts in the baseline specification. With $\sigma = 0.5$ the implications become more noticeable. Modest effects on the government's optimal policy contrast with major changes in equilibrium activity, see figure 6. Since aware households are removed from the infectious pool when information is public but not otherwise, infection risk is smaller in the former case. Unaware households therefore behave less cautiously with public than with private information. Since infections depend on the average activity level when information is private and the relative importance of aware households for that average increases as the epidemic progresses, unaware households become increasingly cautious.

Both in the public and private information cases lockdowns last longer than in the baseline; the lockdown in the private (public) information case is equally (less) strict than in the baseline. The welfare gains from government intervention are lower for the public and higher for the private information cases.

³⁵We maintain the assumption of a negligible death rate.

³⁶Aware households do not bear any costs of the epidemic, neither private nor social, while the remaining population with share $1 - \sigma y$ does, see appendix G.

³⁷COVID-19 symptoms often are private information: Loss of smell and taste is unobserved by third parties and coughing is uninformative about the respiratory disease causing it.




Figure 6: Activity level in equilibrium: Baseline (solid), public information (dashed), and private information (dotted).

5.10 Summary

We draw several conclusions from our quantitative findings. First, the baseline scenario implies an optimal activity reduction during the early phase of the lockdown of roughly 38 percent. When infections are publicly or privately observable this number does not change much even when awareness is high.

Second, changes in fatality rates, the arrival rate of a cure, or regime change have the expected effects on optimal policy, as does the presence of congestion or learning effects; the sensitivity of the policy prescriptions with respect to these parameter changes is small. But three structural factors have major implications for the optimal policy: Much stricter lockdowns are called for when consumption smoothing is less essential or if there is re-infection risk; and much more lenient lockdowns when the activity-infection nexus is linear.

Finally, the model implied welfare gains suggest that government interventions are particularly important in the presence of re-infection risk and when the activity-infection nexus is linear. It is least important with strong consumption smoothing needs. Inverse lockdowns generate a large share of the welfare gains of optimal government intervention.

6 Conclusion

We have developed a flexible economic model of infection dynamics with a single endogenous state variable. Households adjust economic activity for fear of infection but do not internalize static and dynamic externalities. This opens a role for government intervention. We have proved that the optimal policy function is continuous and have established several novel theoretical results whose core intuitions extend to richer frameworks.

First, a lockdown is followed by policies to stimulate activity beyond the privately optimal level. The reason is that, eventually, economic activity generates positive externalities because it drives the economy out of the range with high infections. Second, optimal policy is discontinuous in the deterministic arrival date of a cure: When the arrival date is early then the lockdown is strict; otherwise, it is very loose because the

government is unwilling to accept “lost livelihoods” for as long as would be necessary to save lives. Third, re-infection risk renders the laissez-faire allocation inefficient even in steady state. This calls for permanent interventions when infections become endemic.

Calibrated to the COVID-19 pandemic our baseline specification suggests that starting from mid March 2020, economic activity should immediately have been reduced by 38 percent to yield welfare gains of three and a half percent of lifetime consumption. Robustness along many dimensions contrasts with sensitivity of the policy prescription with respect to the intertemporal elasticity of substitution, activity-infections nexus, and re-infection risk. The government’s ability to impose inverse lockdowns is important and responsible for at least a third of the welfare gains of optimal intervention.

Our workhorse model allows for many other possible extensions. One important avenue for future research concerns additional dimensions of heterogeneity beyond those due to imperfect observability of infection status. A related avenue concerns conflicts of interest and political frictions that affect government interventions.³⁸ Our results on lockdowns and inverse lockdowns also call for a richer analysis of the fiscal consequences of optimal policy.

³⁸[Gonzalez-Eiras and Niepelt \(2022\)](#) analyze the roles of partisanship and career concerns in shaping government responses to the first wave of the COVID-19 epidemic across U.S. states.

References

- Abel, A. B. and Panageas, S. (2020). Optimal management of a pandemic in the short run and the long run, *Working Paper 27742*, NBER, Cambridge, Massachusetts.
- Acemoglu, D., Chernozhukov, V., Werning, I. and Whinston, M. D. (2021). Optimal targeted lockdowns in a multigroup SIR model, *American Economic Review: Insights* **3**(4): 487–502.
- Alvarez, F., Argente, D. and Lippi, F. (2021). A simple planning problem for COVID-19 lockdown, *American Economic Review: Insights* **3**(3): 367–382.
- Atkeson, A. (2020). What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios, *Working Paper 26867*, NBER, Cambridge, Massachusetts.
- Bailey, N. T. J. (1975). *The Mathematical Theory of Infectious Diseases and its Applications*, 2 edn, Hafner Press, New York.
- Bardi, M. and Capuzzo-Dolcetta, I. (1997). *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Springer Science and Business, New York.
- Barles, G. and Souganidis, P. E. (1991). Convergence of approximation schemes for fully nonlinear second order equations, *Asymptotic Analysis* **4**(3): 271–283.
- Bartsch, S. M., Ferguson, M. C., McKinnell, J. A., O’Shea, K. J., Wedlock, P. T., Siegmund, S. S. and Lee, B. Y. (2020). The potential health care costs and resource use associated with COVID-19 in the United States, *Health Affairs* **39**(6): 927–935.
- Bisin, A. and Moro, A. (2022). Learning epidemiology by doing: The empirical implications of a Spatial-SIR model with behavioral responses, *Journal of Urban Economics* **127**: 103368.
- Calvia, A., Gozzi, F., Lippi, F. and Zanco, G. (2023). A simple planning problem for COVID-19 lockdown: A dynamic programming approach, *Economic Theory* pp. <https://doi.org/10.1007/s00199-023-01493-1>.
- Crandall, M. G. and Lions, P.-L. (1983). Viscosity solutions of Hamilton-Jacobi equations, *Transactions of the American Mathematical Society* **277**(1): 1–42.
- Eichenbaum, M. S., Rebelo, S. and Trabandt, M. (2021). The macroeconomics of epidemics, *Review of Financial Studies* **34**(11): 5149–5187.
- Ellison, G. (2020). Implications of heterogeneous SIR models for analyses of COVID-19, *Covid Economics* **53**: 1–32.
- Fajgelbaum, P. D., Khandelwal, A., Kim, W., Mantovani, C. and Schaal, E. (2021). Optimal lockdown in a commuting network, *American Economic Review: Insights* **3**(4): 503–522.

- Farboodi, M., Jarosch, G. and Shimer, R. (2021). Internal and external effects of social distancing in a pandemic, *Journal of Economic Theory* **196**(C): 105293.
- Ferguson, N. M., Laydon, D., Nedjati-Gilani, G. and collaborators (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand, *Report 9*, Imperial College, London.
- Fetzer, T. and Graeber, T. (2020). Does contact tracing work? Quasi-experimental evidence from an Excel error in England, *Working Paper 521*, Centre for Competitive Advantage in the Global Economy, Coventry.
- Garibaldi, P., Pissarides, C. and Moen, E. R. (2020). Static and dynamic inefficiencies in an optimizing model of epidemics, *Discussion Paper 15439*, CEPR, London, UK.
- Gersovitz, M. and Hammer, J. S. (2004). The economical control of infectious diseases, *Economic Journal* **114**: 1–27.
- Giannitsarou, C., Kissler, S. and Toxvaerd, F. (2021). Waning immunity and the second wave: Some projections for SARS-CoV-2, *American Economic Review: Insights* **3**(3): 321–338.
- Gonzalez-Eiras, M. and Niepelt, D. (2020a). On the optimal ‘lockdown’ during an epidemic, *Covid Economics* **7**: 68–87.
- Gonzalez-Eiras, M. and Niepelt, D. (2020b). Optimally controlling an epidemic, *Discussion Paper 15541*, CEPR.
- Gonzalez-Eiras, M. and Niepelt, D. (2020c). Tractable epidemiological models for economic analysis, *Discussion Paper 14791*, CEPR.
- Gonzalez-Eiras, M. and Niepelt, D. (2022). The political economy of early COVID-19 interventions in US states, *Journal of Economic Dynamics and Control* **140**: 104309.
- Hall, R. E., Jones, C. I. and Klenow, P. J. (2020). Trading off consumption and COVID-19 deaths, *Federal Reserve Bank of Minneapolis Quarterly Review* **42**(1): 2–13.
- Hethcote, H. W. (1989). Three basic epidemiological models, in L. Gross, T. G. Hallam and S. A. Levin (eds), *Applied Mathematical Ecology*, Springer, Berlin, pp. 119–144.
- Hethcote, H. W. (2000). The mathematics of infectious diseases, *SIAM Review* **42**(4): 599–653.
- Kaplan, G., Moll, B. and Violante, G. L. (2020). The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the U.S., *Working Paper 27794*, NBER, Cambridge, Massachusetts.
- Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics, *Proceedings of the Royal Society, Series A* **115**(772): 700–721.

- Kolmogorov, A. N. and Fomin, S. V. (1970). *Introductory Real Analysis*, Prentice Hall, New Jersey.
- Menachemi, N., Yiannoutsos, C. T., Dixon, B. E., Duszynski, T. J., Fadel, W. F., Wools-Kaloustian, K. K., Unruh Needleman, N., Box, K., Caine, V., Norwood, C., Weaver, L. and Halverson, P. K. (2020). Population point prevalence of SARS-CoV-2 infection based on a statewide random sample—Indiana, *Morbidity and Mortality Weekly Report*, April 25–29 **69**: 960–964.
- Miclo, L., Spiro, D. and Weibull, J. W. (2022). Optimal epidemic suppression under an ICU constraint: An analytical solution, *Journal of Mathematical Economics* **101**: 102669.
- Piguillem, F. and Shi, L. (2022). Optimal Covid-19 quarantine and testing policies, *Economic Journal* **132**(647): 2534–2562.
- RECOVERY Collaborative Group, Chappell, L., Horby, P., Shen Lim, W., Emberson, J., Mafham, M., Bell, J., Linsell, L., Staplin, N., Brightling, C., Ustianowski, A., Elmahi, E., Prudon, B., Green, C., Felton, T., Chadwick, D., Rege, K., Fegan, C., Faust, S., Jaki, T., Jeffery, K., Montgomery, A., Rowan, K., Juszczak, E., Baillie, J., Haynes, R. and Landray, M. (2020). Dexamethasone in hospitalized patients with Covid-19—preliminary report, *The New England Journal of Medicine* . 10.1056/NEJMoa2021436.
- Toxvaerd, F. (2020). Equilibrium social distancing, *Working Paper 2020/08*, Cambridge Institute for New Economic Thinking, Cambridge.

A Limitations of the One-State-Variable Approach

The long-run share of susceptible households is endogenous in the SIR model while it is exogenous in the logistic framework. This introduces policy tradeoffs that are only present in the SIR model. The question is how relevant these tradeoffs are during the early stage of an epidemic. We answer this question using two approaches. Both imply essentially no relevance.

For the first approach, we choose \bar{y} in the logistic model to best fit the SIR dynamics under the assumption that activity is constant at unity. Then we find the duration D such that reducing activity from 1 to a for this duration and reverting back to unity thereafter lowers z_∞ (the long-run share of the recovered population in the SIR model) by one percent below \bar{y} . The first row of table 3 reports the duration for different activity choices and the second row reports the share of the population that has gone infection after that duration. We conclude from this exercise that in the SIR model policy interventions up to the point at which 3.5 percent of the population have undergone infection, which corresponds to $y = 0.0350$ in the logistic model, have an effect on z_∞ of at most one percent.³⁹ When we double the basic reproduction number then the critical y value rises to 15 percent.

$g(a)$	0.9000	0.8000	0.7000	0.6000	0.5000	0.4500	0.4300	0.4250
D , days	139	151	181	247	448	917	1979	3358
$\iota_D + z_D$	0.2060	0.1044	0.0678	0.0512	0.0402	0.0365	0.0352	0.0349

Table 3: Temporary reductions in $g(a)$ in the SIR model such that $z_\infty = 0.99\bar{y}$.

Second, we contrast the optimal policy in the logistic model with the policy found by Farboodi et al. (2021) who study the government’s problem in the SIR model under the same assumptions as we do in the baseline scenario.⁴⁰ From figure 19 in the online appendix of Farboodi et al. (2021) we recover their optimal policy during the first two years of the epidemic. We fit a fourth-degree polynomial to the activity series and use it as input for a SIR simulation.⁴¹ We associate y in the logistic model with $\iota + z$ obtained in the SIR simulation. This allows us to represent the optimal policy in Farboodi et al. (2021) (as far as their paper reports it, namely for $y \leq 0.0305$) as a function of y , see

³⁹For generality the table reports the results for different values of $g(a)$ rather than a . For $g(a) < 0.4168$ the infected pool shrinks during the intervention. For $g(a) > 0.9768$ the effect on z_∞ is smaller than one percent even if $D \rightarrow \infty$.

⁴⁰Farboodi et al. (2021) stipulate, as we do, a logarithmic benefit and linear cost of economic activity as well as health cost proportional to infections. We refer to the appendix of the paper where the authors analyze the case in which recovered individuals are unaware of their health status, as in our baseline model. We use the parameter values reported in Farboodi et al. (2021).

⁴¹To recover the numerical values we use the software WebPlotDigitizer available at <https://apps.automeris.io/wpd/>.

figure 1. The optimal policy based on the analysis in Farboodi et al. (2021) is practically identical to the one we find in the logistic model.

B Proofs of Lemmas and Propositions

B.1 Proof of Lemma 1

Proof. When $T < \infty$, assumptions 1 and 2 imply $V(y, 0) = U^*$ since there is no reason to reduce activity below the first-best level when infection risk is zero, which is the case when a cure has arrived. Also, $V(y, d) < U^*$ for all $(y, d) \in (0, \bar{y}) \times (0, T]$ since an epidemic involves costs. When $T = \infty$, parallel logic implies $V(0) = V(\bar{y}) = U^*$ and $V(y) < U^*$ for all $y \in (0, \bar{y})$.

Independent of T , assumptions 1 and 2 imply

- i. u and f are continuous and bounded for all $y \in [0, \bar{y}]$, $a \in A$, and A is compact;
- ii. $|f(x, a) - f(y, a)| \leq \beta \bar{y} \max[1, \omega] |x - y|$ for all $x, y \in [0, \bar{y}]$, $a \in A$.

Condition i. and assumptions 1 and 2 ($\rho, \nu \geq 0$; if $T = \infty$ then $\rho + \nu > 0$) imply that V is bounded from above and below. Conditions i. and ii. imply that there exists a unique solution $y(t; \mathbf{a}, y_0)$, $t \in [0, T]$ for each $(y_0, \mathbf{a}) \in [0, \bar{y}] \times \mathcal{A}$ (see section III.5 in Bardi and Capuzzo-Dolcetta (1997), henceforth BCD).

Also independent of T , conditions i. and ii. imply that assumptions A0, A1, A3 and A4 or A4' (depending on whether T is infinite or finite, respectively) in chapter III of BCD are satisfied. Condition A5 in chapter III of BCD is satisfied because P^* is independent of the state. Conditions i. and ii. also imply that condition H1 in chapter II of BCD is satisfied (by remark 3.4 in chapter II). Finally, conditions i. and ii. with assumptions 1 and 2 (finite ω , $u'(\underline{a})$) imply that f and $u - \psi f$ are global Lipschitz continuous in the state space and uniform continuous in the control variable. Lipschitz continuity of f implies that A2 in chapter III of BCD is satisfied.

Under A0, A1, A3 and either A4 or A4' and A5 (depending on whether T is infinite or finite, respectively) the Dynamic Programming Principle holds (propositions 2.5 and 3.2 in section III of BCD, respectively). This implies that the value function V is a viscosity solution of the HJB equation stated in the lemma (propositions 2.8 and 3.5 in chapter III of BCD, respectively).

The viscosity solution is unique. When $T = \infty$ this follows from the fact that any solution of the HJB equation must satisfy the boundary conditions $V(0) = V(\bar{y}) = U^*$, and under H1 two viscosity solutions that are identical on the boundary of the state space must be identical over the entire state space (theorem 3.1 with remarks 3.2 and 3.4 in chapter II of BCD). When $T < \infty$ uniqueness follows under A0–A4' and A5 from theorem 3.7 in chapter III of BCD (noting that A0–A4 imply (2.18) in chapter III of BCD).

When $T < \infty$, A0, A1, A3, A4', A5 and Lipschitz continuity of P^* imply that the value function is Lipschitz continuous (proposition 3.1 (iii) in chapter III of BCD). When $T = \infty$, A0, A1, A3, A4 and Lipschitz continuity in y , uniform continuity in a of $u - \psi f$

imply that V is Hölder continuous with exponent $\min[\frac{\rho+\nu}{\beta\bar{y}\psi\max[1,\omega]}, 1]$ (proposition 2.1 in chapter III of BCD). \square

B.2 Proof of Proposition 1

Proof. Consider first the case of $T < \infty$. Derivations in the text imply that the first-order condition (5) holds at all points in the state space where the value function is differentiable with respect to y . Substituting the first-order condition into the HJB equation therefore implies that

$$(\rho + \nu)V(y, d) = u(a(y, d)) - u'(a(y, d)) \frac{g(a(y, d))}{g'(a(y, d))} - V_d(y, d) + \nu U^* \quad (14)$$

at all such points. From lemma 1, V is Lipschitz continuous. By the Rademacher's theorem this implies that V is differentiable almost everywhere (proposition 1.9 in chapter II of BCD). Condition (14) thus holds almost everywhere in state space.

Consider two points (y_1, d) and (y_2, d) at which V is differentiable. Since V is differentiable almost everywhere we can let $|y_1 - y_2| < \epsilon$ for any $\epsilon > 0$. Continuity of u , u' , g , g' , V_d and V (lemma 1) on the line segment between (y_1, d) and (y_2, d) implies, from equation (14), that the policy function is continuous as well, so $|a(y_1, d) - a(y_2, d)| < w(\epsilon)$ where w denotes a modulus function. Consider a point (y, d) on the line segment and suppose that V is not differentiable at this point. From first-order condition (5) this implies that $a(y, d)$ is not continuous at that point either, contradicting the previous result. We conclude that V is differentiable throughout the state space.

Consider next the case when $T = \infty$. While equation (14) continues to hold (without the V_d term) at all points in the state space where V is differentiable, lemma 1 does not guarantee that V is Lipschitz continuous. However, as shown below, the policy function has bounded variation and since $f(g)$, u as well as the discounting function are of bounded variation the integrand of the value function defined as in (4) is of bounded variation (Kolmogorov and Fomin, 1970, section 32) and therefore integrable. Corollaries 1 and 2 in section 33.2 in Kolmogorov and Fomin (1970) then imply that V is differentiable almost everywhere and an argument parallel to the one when $T < \infty$ implies that V is differentiable throughout the state space.

To establish bounded variation of the policy function suppose to the contrary that there exists no finite partition of the state space such that a is monotone in each subinterval of the partition. Consider a compact subset $[y_a, y_b] \subset [0, \bar{y}]$ of the state space on which a is not of bounded variation and form the partition $[y_0, y_1), [y_1, y_2), \dots, [y_{n-1}, y_n]$ where $y_0 = y_a$, $y_n = y_b$ and $\epsilon \equiv \sup_i \{y_i - y_{i-1}\}_{i=1,2,\dots,n}$ is arbitrarily close to zero. Recall that a is non-monotone in each subinterval i .⁴² Denote by $t(y_i)$ the time at which the economy reaches state y_i and let $\Delta t(\epsilon) = \sup_i \{t(y_i) - t(y_{i-1})\}_{i=1,\dots,n}$. Since $a \geq \underline{a} > 0$ we can make $\Delta t(\epsilon)$ arbitrarily close to zero by appropriately choosing ϵ . Consider an alternative policy \tilde{a} that is constant on each subinterval i and implies the same sequence $\{t(y_i)\}_{i=1,\dots,n}$ as

⁴²Since a is not of bounded variation there exists for each subinterval i a $y \in [y_{i-1}, y_i - \epsilon/3]$ such that $\text{sgn}(a(y + \epsilon/3) - a(y)) = -\text{sgn}(a(y + 2\epsilon/3) - a(y + \epsilon/3))$.

the original policy.⁴³ Policy \tilde{a} has two properties: First, the policies a and \tilde{a} generate the same discounted cost of infections. This follows from the fact that ϵ is arbitrarily close to zero such that discounting within each subinterval i can be disregarded, and because both policies imply the same sequence $\{t(y_i)\}_{i=1,\dots,n}$. Second, \tilde{a} weakly exceeds the average of the original policy on each subinterval because of the convexity of g . Since u is strictly concave and \tilde{a} is in each subinterval both smoother and on average higher than policy a the former policy generates higher discounted utility. We conclude that \tilde{a} dominates a , and thus that the optimal policy must be of bounded variation.

Since independently of T the value function is differentiable, condition (7) applies for any T . Moreover, V approaches U^* as $y \downarrow 0$ such that $\lim_{y \downarrow 0} (\rho + \nu)V(y, T) - \nu U^* = u(a^*)$. Taking the limit for $y \downarrow 0$ in condition (7) thus implies that $u(a) - u'(a)g(a)/g'(a) - u(a^*)$ approaches zero. Since $u(a) - u(a^*) \leq 0$ and $u'(a)g(a)/g'(a) \geq 0$ (with equality if $a = a^*$) it follows that $\lim_{y \downarrow 0} a = a^*$. \square

B.3 Proof of Proposition 2

Proof. Part i. follows because the value function is decreasing in a neighborhood of $y = 0$ (from lemma 1) implying that the static and dynamic externalities both are negative. Part ii. follows from lemma 2, which states that

$$\max_{y \in \mathcal{Y}} V'(y) = \frac{g(a^*)\beta\bar{y}\omega\psi}{\rho + \nu + g(a^*)\beta\bar{y}\omega},$$

and the fact that the total externality is proportional to $V'(y) + \psi(\zeta/n - 1)$. Under the stated condition the total externality therefore eventually (and for $y < \bar{y}$) turns positive and the government imposes an inverse lockdown. Since V is strictly convex in $[y^{\min}, \bar{y}]$ and $y^c \geq y^{\min}$ the total externality switches signs only once, at y^c . \square

B.4 Proof of Proposition 3

Proof. From the Hamiltonian the first-order conditions are

$$\begin{aligned} u'(a(t)) &= g'(a(t))\beta\bar{y}y(t) \left(1 - \left(\frac{y(t)}{\bar{y}}\right)^\omega\right) (\psi - \mu(t)), \\ \dot{\mu}(t) &= g(a(t))\beta\bar{y} \left(1 - (1 + \omega) \left(\frac{y(t)}{\bar{y}}\right)^\omega\right) (\psi - \mu(t)). \end{aligned}$$

Differentiating the first condition with respect to time implies

$$\dot{a}(t) \frac{d}{da(t)} \frac{u'(a(t))}{g'(a(t))} = \beta\bar{y} \left(1 - (1 + \omega) \left(\frac{y(t)}{\bar{y}}\right)^\omega\right) \dot{y}(t) (\psi - \mu(t)) - \beta\bar{y}y(t) \left(1 - \left(\frac{y(t)}{\bar{y}}\right)^\omega\right) \dot{\mu}(t).$$

The right-hand side of this condition collapses to zero when we substitute the second condition and use the law of motion. We conclude that $\dot{a}(t) = 0$ until $t = T$.

⁴³Alternatively, \tilde{a} could be constructed based on a monotone contraction of a on each subinterval of the partition.

Using the functional form assumptions, an optimal allocation thus is characterized by the following two equations in two unknowns, a and $y_T \equiv y(T)$:

$$a = \begin{cases} \frac{1}{1+\beta\bar{y}y_T\left(1-\left(\frac{y_T}{\bar{y}}\right)^\omega\right)^\psi} & \text{if } n = 1 \\ \frac{-1+\sqrt{1+8\beta\bar{y}y_T\left(1-\left(\frac{y_T}{\bar{y}}\right)^\omega\right)^\psi}}{4\beta\bar{y}y_T\left(1-\left(\frac{y_T}{\bar{y}}\right)^\omega\right)^\psi} & \text{if } n = 2 \end{cases}, \quad (15)$$

$$y_T = \frac{\bar{y}}{\left(1 + e^{-a^n\beta\bar{y}\omega T} \left(\left(\frac{\bar{y}}{y(0)}\right)^\omega - 1\right)\right)^{1/\omega}}. \quad (16)$$

Equation (15) represents the government's first-order condition at time T ; it uses the fact that $\mu(T) = 0$. Equation (16) follows from the law of motion, see equation (2). Both conditions can be represented as functions in (y_T, a) space. Specifically, let $a(y_T)$ denote the right-hand side of equation (15) (where n equals either 1 or 2—for ease of notation we do not make the dependence on n explicit) and let

$$A(y_T) \equiv \left[\frac{1}{\beta\bar{y}\omega T} \ln \left(\frac{\left(\frac{\bar{y}}{y(0)}\right)^\omega - 1}{\left(\frac{\bar{y}}{y_T}\right)^\omega - 1} \right) \right]^{1/n}$$

denote the value of a that solves equation (16). A candidate optimal allocation $(a^\diamond, y_T^\diamond)$ satisfies $a^\diamond = a(y_T^\diamond) = A(y_T^\diamond)$.

Note that in the relevant range, function $a(y_T)$ is U-shaped; strictly positive; has a minimum at $y_T = \bar{y}/(1+\omega)^{1/\omega}$; and satisfies $a(0) = 1 = a(\bar{y})$. Function $A(y_T)$ is strictly increasing; satisfies $A(y_0) = 0$; $\lim_{y_T \rightarrow \bar{y}} A(y_T) = \infty$; and an increase in T reduces $A(y_T)$ for any given y_T . This implies the following: (i) For small T , there exists a unique candidate optimal allocation (a, y_T) with $a(y_T) = A(y_T)$ and $y_T \approx y_0$. (ii) For $T \rightarrow \infty$, there exists a unique candidate optimal allocation (a, y_T) with $a(y_T) = A(y_T)$ and $y_T \approx \bar{y}$.

Let

$$\hat{a} = a(\bar{y}/(1+\omega)^{1/\omega}) \equiv \begin{cases} \frac{1}{1+\beta\bar{y}^2\psi\omega/(1+\omega)^{1+1/\omega}} & \text{if } n = 1 \\ \frac{-1+\sqrt{1+8\beta\bar{y}^2\psi\omega/(1+\omega)^{1+1/\omega}}}{4\beta\bar{y}^2\psi\omega/(1+\omega)^{1+1/\omega}} & \text{if } n = 2 \end{cases},$$

$$\hat{T} \equiv \frac{\ln[(\bar{y}/y(0))^\omega - 1]/\omega}{\hat{a}^n\beta\bar{y}\omega} < \infty$$

and note that $0 < \hat{a} < 1$. Suppose that y_0 is sufficiently small, that is, y_0 satisfies both $y_0 < \bar{y}/(1+\omega)^{1/\omega}$ and⁴⁴

$$\frac{\left(1 + \omega^{\frac{1}{\hat{a}^n}} \left(\left(\frac{\bar{y}}{y(0)}\right)^\omega - 1\right)^{1-\frac{1}{\hat{a}^n}}\right)^{1/\omega}}{(1+\omega)^{1/\omega} - \left(1 + \omega^{\frac{1}{\hat{a}^n}} \left(\left(\frac{\bar{y}}{y(0)}\right)^\omega - 1\right)^{1-\frac{1}{\hat{a}^n}}\right)^{1/\omega}} \ln \left(\left[\left(\frac{\bar{y}}{y(0)}\right)^\omega - 1 \right] \frac{1}{\omega} \right) > -\frac{\hat{a}^n\beta\bar{y}^2\psi\omega/(1+\omega)^{1/\omega}}{1 + \ln(\hat{a}) - \hat{a}}. \quad (17)$$

⁴⁴Note that the logarithmic term diverges as $y_0 \rightarrow 0$.

Starting from $T \approx 0$ and thus, $y_T \approx y_0$, consider a continuous increase in T . This scales the $A(y_T)$ schedule down such that the intersection of the $a(y_T)$ and $A(y_T)$ schedules shifts to the right and down in (y_T, a) space. When T reaches \hat{T} , the intersection reaches the point $(y_T, a) = (\bar{y}/(1 + \omega)^{1/\omega}, \hat{a})$. The welfare generated under the policy $a = \hat{a}$ given $T = \hat{T}$ equals $(1 + \ln(\hat{a}) - \hat{a})\hat{T} - \psi\bar{y}/(1 + \omega)^{1/\omega}$. In contrast, welfare under the policy $a = a^*$ given $T = \hat{T}$ equals

$$\begin{aligned} & (1 + \ln(a^*) - a^*)\hat{T} - \psi \frac{\bar{y}}{\left(1 + e^{-\beta\bar{y}\omega\hat{T}} \left(\left(\frac{\bar{y}}{y(0)}\right)^\omega - 1\right)\right)^{1/\omega}} \\ = & 0 - \psi \frac{\bar{y}}{\left(1 + \omega^{\frac{1}{\hat{a}^n}} \left(\left(\frac{\bar{y}}{y(0)}\right)^\omega - 1\right)^{1 - \frac{1}{\hat{a}^n}}\right)^{1/\omega}}. \end{aligned}$$

Under condition (17), this last term strictly exceeds $(1 + \ln(\hat{a}) - \hat{a})\hat{T} - \psi\bar{y}/(1 + \omega)^{1/\omega}$. Accordingly, $a = \hat{a}$ is suboptimal when $T = \hat{T}$; that is, when $T = \hat{T}$ there exists an activity level $a \neq \hat{a}$ which also solves the equation system (15)–(16) and dominates the choice $a = \hat{a}$.

We have established that the system (15)–(16) has a unique solution for $T \approx 0$. We have also established that the system (15)–(16) has multiple solutions for $T = \hat{T}$ and that the optimal solution is not of the type that arises when $T \approx 0$. It follows that there exists some $T^\circ < \hat{T}$ at which the optimal solution ceases to be of the type that arises when $T \approx 0$. That is, a is discontinuous at $T = T^\circ$. \square

B.5 Proof of Proposition 4

Proof. We prove the first part of the proposition by contradiction. Suppose that $(a(\infty), y(\infty))$ is a steady state both in equilibrium and in the planner's program. Note that the planner's and the equilibrium first-order conditions resemble conditions (5) and (8) in the baseline. Evaluating the conditions at the candidate steady state and using the fact that steady-state gross infection flows are $\gamma y(\infty)$ yields the following expressions for the static and dynamic externalities, respectively:

$$\begin{aligned} \mathcal{S}(a(\infty), y(\infty)) &= \psi\gamma y(\infty) \left(\frac{\zeta}{a(\infty)} - \frac{g'(a(\infty))}{g(a(\infty))} \right), \\ \mathcal{D}(a(\infty), y(\infty)) &= \gamma y(\infty) \frac{g'(a(\infty))}{g(a(\infty))} V'(y(\infty)). \end{aligned}$$

Generically, the sum of these two externalities differs from zero. This implies the contradiction.

Turning to the second part of the proposition, note that the system of equations

characterizing the equilibrium steady state is given by

$$\begin{aligned} y(\infty) &= \bar{y} \left(1 - \frac{\gamma}{\beta \bar{y} a^n(\infty)} \right)^{1/\omega}, \\ a(\infty) &= \begin{cases} \frac{1}{1 + \beta \bar{y} y(\infty) \left(1 - \left(\frac{y(\infty)}{\bar{y}} \right)^\omega \right) \psi \zeta} & \text{if } n = 1 \\ \frac{-1 + \sqrt{1 + 4\beta \bar{y} y(\infty) \left(1 - \left(\frac{y(\infty)}{\bar{y}} \right)^\omega \right) \psi \zeta}}{2\beta \bar{y} y(\infty) \left(1 - \left(\frac{y(\infty)}{\bar{y}} \right)^\omega \right) \psi \zeta} & \text{if } n = 2 \end{cases}. \end{aligned}$$

Eliminating $y(\infty)$ yields generalized polynomials in $a(\infty)$ which have a unique solution when $\gamma = 0$. Totally differentiating the polynomials yields, both for $n = 1$ or $n = 2$, $\lim_{\gamma \rightarrow 0} da(\infty)/d\gamma = -\psi \bar{y} \zeta$.

The steady state implemented by the government is characterized by three equations in the three unknowns $a(\infty)$, $y(\infty)$, and $V'(y(\infty))$, namely

$$\begin{aligned} y(\infty) &= \bar{y} \left(1 - \frac{\gamma}{\beta \bar{y} a^n(\infty)} \right)^{1/\omega}, \\ a(\infty) &= \begin{cases} \frac{1}{1 + \beta \bar{y} y(\infty) \left(1 - \left(\frac{y(\infty)}{\bar{y}} \right)^\omega \right) (\psi - V'(y(\infty)))} & \text{if } n = 1 \\ \frac{-1 + \sqrt{1 + 8\beta \bar{y} y(\infty) \left(1 - \left(\frac{y(\infty)}{\bar{y}} \right)^\omega \right) (\psi - V'(y(\infty)))}}{4\beta \bar{y} y(\infty) \left(1 - \left(\frac{y(\infty)}{\bar{y}} \right)^\omega \right) (\psi - V'(y(\infty)))} & \text{if } n = 2 \end{cases}, \\ (\rho + \nu)V'(y(\infty)) &= -a^n(\infty)\beta \bar{y} \left(1 - (1 + \omega) \left(\frac{y(\infty)}{\bar{y}} \right)^\omega \right) (\psi - V'(y(\infty))) - \gamma V'(y(\infty)). \end{aligned}$$

Eliminating $y(\infty)$ and $V'(y(\infty))$ yields generalized polynomials in $a(\infty)$ which have a unique solution when $\gamma = 0$. Totally differentiating the polynomials yields, both for $n = 1$ or $n = 2$, $\lim_{\gamma \rightarrow 0} da(\infty)/d\gamma = -n\psi(\rho + \nu)\bar{y}/(\rho + \nu + \beta \bar{y}\omega)$.

Since for $\gamma = 0$ the steady-state activity levels in equilibrium and under the optimal policy are identical we conclude that for small values of γ the government chooses a higher steady-state activity level (and a higher $y(\infty)$ ⁴⁵) than in equilibrium if and only if $-\psi \bar{y} \zeta < -n\psi(\rho + \nu)\bar{y}/(\rho + \nu + \beta \bar{y}\omega)$ or

$$(\rho + \nu)(n - \zeta) < \zeta \beta \bar{y} \omega.$$

□

⁴⁵Recall that in steady state $y(\infty)$ is given by the expression in condition (11).

C Relation to SIR and SIS Models

The law of motion (1) nests a range of well known epidemiological models, augmented by an effect of economic activity on the infection rate. For example, under the restriction $\omega = \bar{y} = 1$ it corresponds to the SI model (Bailey, 1975) in which the number of transitions from the pre to the post group first increases and then decreases as $y(t)$ moves from near zero towards unity.

The law of motion also corresponds to a special case of the canonical SIR model (Kermack and McKendrick, 1927) and the modified SIR model (Bailey, 1975). In their general form, both SIR models characterize the evolution of three population groups—susceptible, currently infected, and removed households—and therefore contain two endogenous state variables (the shares of two of the three groups). Susceptible households are infected by currently infected households; they remain infectious for a random time span; and eventually they join the pool of removed (recovered or deceased) households.

These dynamics reduce to the law of motion (1) with $\omega = 1$ when the distinction between currently infected and removed households is blurred and they are combined into a single group of post households.⁴⁶ As explained in the main text this does not undermine the model’s capacity to represent societal costs of infection or death.

While the canonical SIR model (Kermack and McKendrick, 1927) and the modified SIR model (Bailey, 1975) yield similar predictions for transition dynamics they differ with respect to their implications for the long-run share of the population that never gets infected. In the canonical SIR model this long-run share and the “herd-immunity” level are endogenous while in the modified SIR model the share equals zero. A hybrid model augments the modified SIR model with an additional parameter that allows to regulate the long-run population shares (Gonzalez-Eiras and Niepelt, 2020c). In the law of motion (1) the parameter \bar{y} plays a similar role: A lower \bar{y} implies a larger share of the population that never gets infected.

The law of motion (1) also is closely related to the SIS model in which households, once infected, randomly recover and return to the susceptible pool (rather than the removed pool as in SIR models) because infection does not confer immunity (Hethcote, 1989).⁴⁷ Our framework differs from the SIS model insofar as we represent infections in terms of *flows*, $\dot{y}(t)$ implied by (1), while in the SIS model the key endogenous variable represents the *stock* of the infected population.

D Calibration Strategy and Numerical Methods

D.1 Law of Motion for Infections

In this appendix, we describe how we use information about parameter values in the canonical SIR model to deduce parameter values for the law of motion (1).

⁴⁶That is, when the transition rate from currently infected to removed in the SIR model equals zero.

⁴⁷The SIS model has $\bar{y} = 1$.

D.1.1 Canonical SIR Model

The canonical SIR model due to [Kermack and McKendrick \(1927\)](#) specifies laws of motion for the population shares of three groups: the “susceptible,” the “infected” or “infectives,” and the “removed.” Their respective population shares at time $t \geq 0$ are denoted by $x(t)$, $\iota(t)$, and $z(t)$, respectively, where $x(t) + \iota(t) + z(t) = 1$.⁴⁸ We normalize the mass of the total population at time $t = 0$ to unity.

At time $t = 0$ the population consists of $x(0)$ susceptible persons and a few infected persons, $\iota(0)$. There are no removed persons at this time, $z(0) = 0$. In each instant after time $t = 0$, infected persons transmit the disease to members of the susceptible group and a share of the infected either dies or recovers and develops immunity. Formally,

$$\dot{x}(t) = -b(t)x(t)\iota(t), \quad (18)$$

$$\dot{\iota}(t) = -\dot{x}(t) - (c^d + c^r)\iota(t), \quad (19)$$

$$\dot{z}(t) = (c^d + c^r)\iota(t). \quad (20)$$

Here, $b(t)$ denotes a possibly time-varying infection rate. The extent to which susceptible persons are infected depends on their number, $x(t)$; the infection rate, $b(t)$; and the population share of infected persons. The number of infected persons increases one-to-one with the susceptible persons that get infected, while a share $c \equiv c^d + c^r$ of the infected population dies or recovers; the coefficients c^d and c^r parameterize the flow into death and recovery, respectively.

Consider the case where $b(t)$ is constant at value b . Inspection of equations (18) and (19) reveals that for $bx(0) > c$ the share of infected persons increases until it reaches a maximum when $x(t) = c/b$; thereafter, the share declines. Intuitively, when $x(0)$ falls short of c/b (the “herd immunity level”) then there are fewer new infections of susceptible persons than outflows from the infected pool due to recoveries and death. As is well known (e.g., Theorem 2.1 in [Hethcote, 2000](#)), $x(\infty)$ falls short of the herd immunity level unless $x(0) = c/b = x(\infty)$ and $\iota(0) = 0$.⁴⁹

In the SIR-S model a share γ of the removed population loses immunity and moves back to the susceptible pool. Accordingly, the dynamic system is given by

$$\dot{x}(t) = -b(t)x(t)\iota(t) + \gamma z(t),$$

$$\dot{\iota}(t) = b(t)x(t)\iota(t) - c\iota(t),$$

$$\dot{z}(t) = c\iota(t) - \gamma z(t).$$

In steady state this reduces to

$$\gamma z = bx\iota = c\iota.$$

⁴⁸We follow the notation introduced by [Kermack and McKendrick \(1927\)](#) except for denoting the share of infected by ι rather than y .

⁴⁹Note also, from equation (19), that at the beginning of an epidemic with $x(t) \approx 1$ and $z(t) \approx 0$, b approximately equals the growth rate of the number of persons who are or were infected, $\frac{\dot{\iota}(t) + \dot{z}(t)}{\iota(t) + z(t)} = b \frac{x(t)\iota(t)}{\iota(t) + z(t)} \approx b$.

Calibration We measure time in days and use information about the spread of COVID-19 in the United States to calibrate the model. We associate $t = 0$ with 16 March 2020, the date at which the median U.S. state closed schools and public health authorities considered further restrictions.⁵⁰

Following Atkeson (2020) and the sources cited therein we assume that the flow rate from the infected to the removed population equals $c = 1/18$, corresponding to an exponentially distributed infection duration that averages 18 days.⁵¹ To calibrate b we rely on information in Ferguson et al. (2020) who argue that the “basic reproduction number” $\mathcal{R}_0 = b/c$ for COVID-19 equals approximately 2.4. This implies $b = 0.1333$.

By March 16, 2020 the U.S. reported a total of 91 COVID-19 deaths out of a population of 328 million.⁵² With an infection fatality rate of 0.58 percent (Menachemi et al., 2020) this implies $z(0) = 91/(0.58\% \cdot 328 \cdot 10^6) = 0.4783 \cdot 10^{-4}$.⁵³ We use the following well-known result (e.g., Theorem 2.1 in Hethcote, 2000):

Proposition 5. Let $s < 0$ denote a start date. In the canonical SIR model with $\iota(s) > 0$ and $b(t) = b > c/x(s)$,

$$1 - z(0) = x(0) + \iota(0) = x(s) + \iota(s) + \frac{c}{b} \ln \left(\frac{x(0)}{x(s)} \right).$$

The long-run share of the susceptible population, $x(\infty)$, solves the equation

$$x(\infty) = x(0) + \iota(0) + \frac{c}{b} \ln \left(\frac{x(\infty)}{x(0)} \right).$$

Since the number of infected or removed persons was negligible before mid March the first set of equalities implies

$$1 - z(0) = 1 + \frac{1}{2.4} \ln \left(\frac{x(0)}{1} \right) \Rightarrow x(0) = 1 - 0.1148 \cdot 10^{-3}, \quad \iota(0) = 0.6696 \cdot 10^{-4}.$$

The second condition implies $x(\infty) = 0.1214$.

D.1.2 Generalized Logistic Model

Given $y(0) = \iota(0) + z(0) = 0.6696 \cdot 10^{-4} + 0.4783 \cdot 10^{-4} = 0.1148 \cdot 10^{-3}$ and $\bar{y} = 1 - x(\infty) = 0.8786$, we choose β and ω in equation (1) (subject to $g(a) = 1$) to best fit the implied time series for $\dot{y}(t)$ to the path for $\iota(t)$ in the calibrated SIR model (allowing for an arbitrary factor of proportionality). This yields $\beta = 0.8346 \cdot 10^{-1}$ and $\omega = 0.6662$. Figure 7 illustrates the very close parallels between the predictions of the canonical SIR model (in blue) and the generalized logistic model (in black).

⁵⁰The median state closed restaurants and imposed restrictions on public gatherings on 17 March 2020. Puerto Rico imposed a stay-at-home order on 15 March 2020. See <https://web.csg.org/covid19/executive-orders/>.

⁵¹Note that $\int_0^\infty ce^{-ct} dt = 1/c$.

⁵²See <https://github.com/nytimes/covid-19-data/blob/master/us.csv>.

⁵³See Farboodi et al. (2021) for similar calculations.




Figure 7: Dynamics in the canonical SIR model (blue) and in the generalized logistic model (black). SIR model: $x(t)$ (dotted), $\iota(t)$ (solid), and $z(t)$ (dashed). Generalized logistic model: $\dot{y}(t)$ (solid, scaled), and $y(t)$ (dashed).

D.2 Costs of Infection

In this appendix, we discuss the calibration of the parameters representing private and social costs of infection. Recall that $\xi = \zeta \frac{a_i}{a} + (1 - \zeta)$. We calibrate ζ based on U.S. estimates of hospitalization costs and the value of life by [Bartsch et al. \(2020\)](#) and [Hall et al. \(2020\)](#), respectively. [Bartsch et al. \(2020\)](#) estimate direct medical costs including follow up expenses (over a year) of \$1.25 trillion under the assumption that eighty percent of the U.S. population are infected. We scale this number reflecting our modified specification of $\bar{y} = 1 - x(\infty)$ (0.8786 rather than 0.8000); this yields (unconditional) costs per capita of about \$4,185. [Hall et al. \(2020\)](#) assess the value of life at \$270,000 per year. With an average remaining life expectancy of 14.5 years every life lost to COVID-19 thus costs \$3,915,000. Conditional on the infection fatality rate of 0.58 percent ([Menachemi et al., 2020](#)) and $\bar{y} = 0.8786$ this implies (unconditional) costs due to COVID-19 deaths of \$19,950 per capita. Under the assumption that individuals fully internalize mortality risk but not marginal social medical costs we conclude that $\zeta = 19,950 / (19,950 + 4,185) = 0.8266$.

To calibrate ψ based on the dollar amount \$19,950 + \$4,185 we follow the approach of [Hall et al. \(2020\)](#) who quantify willingness to pay for reduced mortality risk. Under our baseline assumptions of an infection risk of 0.8786, an infection fatality rate of 0.58 percent, and relative risk aversion of one we find that an individual would sacrifice a share $1 - \phi = 0.3581$ of consumption to eliminate COVID-19 related mortality risk (neglecting

other costs).⁵⁴ In the model the annual utility costs of sacrificing this share then equal⁵⁵

$$365 \cdot \{(1 + \ln(a^*) - a^*) - (1 + \ln(a^*\phi) - a^*)\} = -365 \ln(\phi).$$

Since without reductions in activity almost all infections occur during the first year ($\int_0^{365} \dot{y}(t) dt \approx \bar{y}$ where the trajectory $y(t)$ is evaluated at $a = 1$) we conclude that the social cost parameter reflecting mortality risk, $\hat{\psi}$, equals $\hat{\psi} = -365 \ln(\phi)/\bar{y} = 184.2$. Adding medical costs we arrive at an estimate for $\psi = \hat{\psi}/\zeta$ of 222.8.⁵⁶

Endogenous Costs When ψ is a function of the state we need to modify the calibration. Throughout, we evaluate the trajectory $y(t)$ at $a = 1$. To capture congestion effects we replace ψ in (3) by $\psi^f y(t)(1 - (\frac{y(t)}{\bar{y}})^\omega)$ where $\psi^f > 0$. Using

$$\int_0^\infty \dot{y}(t)y(t) \left(1 - \left(\frac{y(t)}{\bar{y}}\right)^\omega\right) dt = \int_0^{\bar{y}} y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) dy = \frac{\omega \bar{y}^2}{2(2 + \omega)}$$

and under the maintained assumption that almost all infections occur during the first year we conclude that $\psi^f = \psi 2(2 + \omega)/(\omega \bar{y})$.

Suppose next that we replace ψ in (3) by $\psi(1 - \ell y(t)/\bar{y})$, representing learning-by-doing. Consistent with evidence we assume that the learning effects let unit costs drop by a third over the course of the epidemic: $\ell = 1/3$.⁵⁷

D.3 Numerical Methods

We use Mathematica to numerically solve for the value function and associated policy function. The domain is discretized using fourth-order finite difference methods to approximate the continuous HJB equation. The value function that solves the discretized HJB equation defined over the continuous state space converges to the value function of the original continuous HJB under our model assumptions (see [Bardi and Capuzzo-Dolcetta \(1997, theorem 1.1, section VI\)](#)). The method of lines is used for time integration

⁵⁴Let ℓ denote expected costs due to loss of life (which we calibrate to \$19,950, see above) relative to average consumption (which [Hall et al. \(2020\)](#) set to \$45,000). This ratio, which represents the sacrifice ratio of an agent with linear preferences, equals 0.4433. [Hall et al. \(2020\)](#) derive ϕ from ℓ based on the relationship $\phi = [1 + (r - 1)\ell]^{1/(1-r)}$ where r denotes the coefficient of relative risk aversion. For $r = 0$ this yields $1 - \phi = \ell$; for $r \rightarrow 1$ (logarithmic utility), $1 - \phi \rightarrow 1 - \exp(-\ell)$ or roughly 0.3581. For $r = 2$ we find $1 - \phi = 0.3071$. [Hall et al. \(2020\)](#) assume CRRA preferences, a coefficient of relative risk aversion of two, and mortality risk of 0.44 percent per year. They report that the share of consumption that an individual would sacrifice to eliminate COVID-19 related mortality risk equals 28 percent.

⁵⁵We neglect time discounting as do [Hall et al. \(2020\)](#). Note that only the benefit of economic activity (“consumption”) not the cost associated with it (“labor supply”) is reduced by the fraction $1 - \phi$.

⁵⁶The calibration of ψ is independent of our preference assumption. If we stipulate $r = 2$ rather than $r = 1$ (see footnote 54) then the annual utility costs of sacrificing the consumption share $1 - \phi$ equals $365 \cdot \{-(a^*)^{-1} + (a^*\phi)^{-1}\} = -365(1 - \phi^{-1})$. But $\ln(1 - 0.3581) \approx 1 - (1 - 0.3071)^{-1}$, so the different value for r does not materially affect ψ .

⁵⁷[RECOVERY Collaborative Group et al. \(2020\)](#) conducted a randomized trial in the UK finding that among 2104 patients that received dexamethasone mortality was reduced by one-third in patients receiving mechanical ventilation and by one-fifth in patients receiving oxygen but not mechanical ventilation.

and standard numerical methods for solving ODEs, such as the Runge-Kutta method, are used to solve for the discretized HJB. [Barles and Souganidis \(1991\)](#) show that the finite difference method satisfies the monotonicity, consistency and stability conditions that guarantee convergence of the approximation to the unique viscosity solution of the HJB equation.

As a robustness check we solve the HJB equation using the Initial Value Problem Differential-Algebraic Equations (IDA) method. This is a robust and efficient method designed for differential-algebraic equations that exploits the algebraic constraints in ODEs. The results we obtain are essentially identical.

E Viscosity Solutions of Nonlinear Partial Differential Equations

A generic HJB equation takes the form $v_d(y, d) + F(y, v(y, d), D_y v(y, d)) = 0$ where $v : \Omega \times [0, T] \rightarrow \mathbb{R}$ is a continuous value function; $\Omega \subseteq \mathbb{R}^n$ and $[0, T]$ denote the endogenous and exogenous state spaces, respectively; $D_y v(y, d)$ denotes the gradient with respect to the endogenous state; and F is smooth. [Crandall and Lions \(1983\)](#) introduce the notion of viscosity solution of a partial differential equation such as this HJB equation; for a textbook treatment see, e.g., [Bardi and Capuzzo-Dolcetta \(1997\)](#).

For all $(y, d) \in \Omega \times [0, T]$ define the superdifferential and subdifferential, respectively, of v as the following sets:

$$\begin{aligned} D^+ v(y, d) &= \left\{ p \in \mathbb{R}^n : \limsup_{x \rightarrow y} \frac{v(x, d) - v(y, d) - p \cdot (x - y)}{|x - y|} \leq 0 \right\}, \\ D^- v(y, d) &= \left\{ q \in \mathbb{R}^n : \liminf_{x \rightarrow y} \frac{v(x, d) - v(y, d) - q \cdot (x - y)}{|x - y|} \geq 0 \right\}. \end{aligned}$$

Note that if both $D^+ v(y, d)$ and $D^- v(y, d)$ are non-empty then $D^+ v(y, d) = D^- v(y, d)$ and v is differentiable at (y, d) .

A continuous function w is a viscosity subsolution of the HJB equation if

$$w_d(y, d) + F(y, w(y, d), p) \leq 0 \quad \forall p \in D^+ w(y, d), \quad \forall (y, d) \in \Omega \times [0, T].$$

It is a viscosity supersolution if

$$w_d(y, d) + F(y, w(y, d), q) \geq 0 \quad \forall q \in D^- w(y, d), \quad \forall (y, d) \in \Omega \times [0, T].$$

If w is a viscosity subsolution and supersolution of the HJB equation then w is a viscosity solution of the HJB equation.

F Other Properties of the Value Function

Lemma 2. Under assumptions [1](#) and [2](#) and if $T = \infty$, V has a unique minimum at $y^{\min} < \bar{y}/(1 + \omega)^{1/\omega}$. Parameter changes that imply a higher (lower) y^{\min} also imply less

(more) pronounced convexity of V at y^{\min} . Moreover,

$$\max_{y \in [y^{\min}, \bar{y}]} V'(y) = V'(\bar{y}) = \frac{g(a^*)\beta\bar{y}\omega\psi}{\rho + \nu + g(a^*)\beta\bar{y}\omega} < \psi$$

and V is strictly convex over the domain $[y^{\min}, \bar{y}]$.

Proof. For the proof and henceforth we assume that V is twice differentiable. From the government's HJB equation, the envelope condition with $T = \infty$ reads

$$(\rho + \nu)V'(y) = -g(a(y))\beta\bar{y} \times \left[\left(1 - (1 + \omega) \left(\frac{y}{\bar{y}} \right)^\omega \right) (\psi - V'(y)) - y \left(1 - \left(\frac{y}{\bar{y}} \right)^\omega \right) V''(y) \right]. \quad (21)$$

Let \hat{y} denote a point where V reaches a local minimum or maximum. Evaluated at \hat{y} the envelope condition reduces to

$$\left(1 - (1 + \omega) \left(\frac{\hat{y}}{\bar{y}} \right)^\omega \right) \psi = \hat{y} \left(1 - \left(\frac{\hat{y}}{\bar{y}} \right)^\omega \right) V''(\hat{y}). \quad (22)$$

Note that any extremum is locally unique since $V''(\hat{y}) = 0$ would only be consistent with $\hat{y} = \bar{y}/(1 + \omega)^{1/\omega}$.

Uniqueness of y^{\min} To see that V has a unique minimum suppose to the contrary that there exist multiple local minima. Consider two neighboring minima at, say, y^a and y^c with $y^a < y^c$. Then there must exist a local maximum at some y^b with $y^a < y^b < y^c$. Since $V''(y^a) > 0$, $V''(y^b) < 0$, and $V''(y^c) > 0$ the right-hand side of condition (22) evaluated at y^a , y^b , and y^c is strictly positive, negative, and positive, respectively. The sign of the left-hand side of condition (22) evaluated at the same points cannot alternate in this way. We have thus arrived at a contradiction which proves that V has a unique minimum, y^{\min} .

Upper Bound on y^{\min} To derive a contradiction suppose first that $y^{\min} > \bar{y}/(1 + \omega)^{1/\omega}$. Since the minimum is unique this implies $V'(\bar{y}/(1 + \omega)^{1/\omega}) < 0$. From the envelope condition (21),

$$(\rho + \nu)V'(\bar{y}/(1 + \omega)^{1/\omega}) = g(a(\bar{y}/(1 + \omega)^{1/\omega}))\beta\bar{y}^2 \frac{\omega}{(1 + \omega)^{1+1/\omega}} V''(\bar{y}/(1 + \omega)^{1/\omega})$$

and thus $V''(\bar{y}/(1 + \omega)^{1/\omega}) < 0$. Since by assumption $y^{\min} > \bar{y}/(1 + \omega)^{1/\omega}$ there must exist an inflection point, say y^i , with $\bar{y}/(1 + \omega)^{1/\omega} < y^i < y^{\min}$, $V'(y^i) < 0$, and $V''(y^i) = 0$. But evaluated at y^i the envelope condition implies

$$\underbrace{(\rho + \nu)V'(y^i)}_{<0} = - \underbrace{g(a(y^i))\beta\bar{y}}_{>0} \underbrace{\left(1 - (1 + \omega) \left(\frac{y^i}{\bar{y}} \right)^\omega \right)}_{<0} \underbrace{(\psi - V'(y^i))}_{>0},$$

which yields a contradiction. We conclude that $y^{\min} \leq \bar{y}/(1 + \omega)^{1/\omega}$.

In fact, $y^{\min} < \bar{y}/(1 + \omega)^{1/\omega}$; for if y^{\min} equalled $\bar{y}/(1 + \omega)^{1/\omega}$ then the minimum could not be unique since condition (22) would imply $V''(\bar{y}/(1 + \omega)^{1/\omega}) = 0$. We conclude that $y^{\min} < \bar{y}/(1 + \omega)^{1/\omega}$.

Effect of Parameters on Convexity of V at y^{\min} From equation (22), $V''(y^{\min}) = \left(1 - (1 + \omega) \left(\frac{y^{\min}}{\bar{y}}\right)^\omega\right) \psi / y^{\min} / \left(1 - \left(\frac{y^{\min}}{\bar{y}}\right)^\omega\right)$. This is strictly decreasing in y^{\min} . The result then follows.

Maximum Slope of V Let $\mathcal{Y} = [y^{\min}, \bar{y}]$ and suppose first that there exists no $y \in \mathcal{Y}$ such that $V''(y) = 0$ (no inflection point on the domain \mathcal{Y}). Since $V(y)$ has a unique global minimum at y^{\min} , $V(y)$ is strictly increasing and convex in this case for all $y \in \mathcal{Y} \setminus y^{\min}$, and thus $\max_{y \in \mathcal{Y}} V'(y) = V'(\bar{y})$. Suppose next that there exists some $y \in \mathcal{Y}$ such that $V''(y) = 0$ (at least one inflection point on the domain \mathcal{Y}). Let $\mathcal{Y}^i \subset \mathcal{Y}$ denote the set of inflection points and let $y^i = \arg \max_{y \in \mathcal{Y}^i} V'(y)$. Then, $\max_{y \in \mathcal{Y}} V'(y) = \max[V'(y^i), V'(\bar{y})]$. From the envelope condition (21) and the fact that $V''(y^i) = 0$, $(\rho + \nu)V'(y^i) = -g(a(y^i))\beta\bar{y} \left(1 - (1 + \omega) \left(\frac{y^i}{\bar{y}}\right)^\omega\right) (\psi - V'(y^i))$ or

$$V'(y^i) = \frac{g(a(y^i))\beta\bar{y} \left((1 + \omega) \left(\frac{y^i}{\bar{y}}\right)^\omega - 1 \right) \psi}{(\rho + \nu + g(a(y^i))\beta\bar{y} \left((1 + \omega) \left(\frac{y^i}{\bar{y}}\right)^\omega - 1 \right))} \leq \frac{g(a^*)\beta\bar{y}\omega\psi}{\rho + \nu + g(a^*)\beta\bar{y}\omega} = V'(\bar{y}),$$

where the weak inequality follows from $(1 + \omega) \left(\frac{y^i}{\bar{y}}\right)^\omega - 1 \leq \omega$, $a(y^i) \leq a^*$, and the fact that g is increasing; and the equality on the right-hand side follows directly from condition (21). Accordingly, $\max_{y \in \mathcal{Y}} V'(y) = V'(\bar{y})$. We conclude that in either case $\max_{y \in \mathcal{Y}} V'(y) = V'(\bar{y})$.

Strict Convexity of V on the Domain \mathcal{Y} From the previous results $V(y)$ is weakly convex in a neighborhood of \bar{y} and strictly convex at y^{\min} .

[1] Suppose first that there do not exist open intervals on the domain \mathcal{Y} such that $V''(y) = 0$ for all points in the interval (however, there may exist points $y \in \mathcal{Y}$ with $V''(y) = 0$). Given the convexity of V at \bar{y} and y^{\min} , the number of inflection points on the domain \mathcal{Y} at which $V'(y)$ changes signs must be even (including equal to zero). To arrive at a contradiction suppose that the number is strictly positive and consider the smallest two neighboring inflection points on the domain \mathcal{Y} , say, y^a and y^b with $y^a < y^b$. Given the specified assumptions, $0 < V'(y^a) < \psi$ (using the above argument on the maximum slope of V), $V''(y^a) = V''(y^b) = 0$, and $V''(y) < 0 < V'(y)$ for all $y \in (y^a, y^b)$. Note that $y^a > \bar{y}/(1 + \omega)^{1/\omega}$ since the envelope condition (21) and $V''(y^a) = 0$ imply $(\rho + \nu)V'(y^a) = -g(a(y^a))\beta\bar{y} \left(1 - (1 + \omega) \left(\frac{y^a}{\bar{y}}\right)^\omega\right) (\psi - V'(y^a))$ such that $0 < V'(y^a) < \psi$ requires $1 - (1 + \omega) \left(\frac{y^a}{\bar{y}}\right)^\omega < 0$. Accordingly, $1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega < 0$ for all $y \in (y^a, y^b)$.

Rewriting the envelope condition (21) as

$$g(a(y)) = -\frac{(\rho + \nu)V'(y)}{\beta\bar{y} \left[\left(1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V'(y)) - y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) V''(y) \right]} \quad (23)$$

and differentiating with respect to y gives, for all $y \in (y^a, y^b)$,

$$\begin{aligned}
g'(a(y))a'(y) &= -\frac{(\rho + \nu)V''(y)}{\beta\bar{y} \left[\left(1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V'(y)) - y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) V''(y) \right]} \\
&\quad + \frac{(\rho + \nu)V'(y)D}{\beta\bar{y} \left[\left(1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V'(y)) - y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) V''(y) \right]^2} \\
&= \underbrace{\frac{V''(y)}{V'(y)}g(a(y))}_{<0} + \underbrace{(\rho + \nu)V'(y)}_{>0} \frac{-\{>0\} - 2\{>0\} - \{>0\}V'''(y)}{\{>0\}},
\end{aligned}$$

where $D \equiv (V'(y) - \psi)\omega(1 + \omega)\frac{y^{\omega-1}}{\bar{y}^\omega} - 2 \left(1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega\right) V''(y) - y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) V'''(y)$ and $\{>0\}$ denotes strictly positive terms (which might differ from each other). Since V is strictly concave for $y \in (y^a, y^b)$ there exists some nonempty $\mathcal{Z} \subseteq (y^a, y^b)$ with $V'''(y) > 0$ for all $y \in \mathcal{Z}$. Since g is increasing the preceding equality implies that $a'(y) < 0$ for all $y \in \mathcal{Z}$.

The government's first-order condition implies

$$\frac{u'(a(y))}{g'(a(y))} = \beta\bar{y}y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V'(y)). \quad (24)$$

Differentiating the right-hand side with respect to y yields

$$\beta\bar{y} \left[\left(1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V'(y)) - y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) V''(y) \right],$$

which is strictly negative for all $y \in (y^a, y^b)$ because of condition (23) and the fact that $g(a(y)) > 0$. Accordingly, the right-hand side of equation (24) is strictly decreasing in y , and so must be the left-hand side. This requires $a'(y) > 0$ for all $y \in (y^a, y^b)$ since g is weakly convex (assumption 1) and u strictly concave (assumption 2). In particular, $a'(y) > 0$ for all $y \in \mathcal{Z}$. We have thus arrived at a contradiction and conclude that there exist no inflection points on the domain \mathcal{Y} .

[2] Suppose next that there do exist open intervals on the domain \mathcal{Y} such that $V''(y) = 0$ for all points in the interval (there may also exist inflection points). Consider such an open interval, say, (y^a, y^b) . Then, $V''(y) = 0$ for all $y \in (y^a, y^b)$. Differentiating the envelope condition with respect to y gives, for all $y \in (y^a, y^b)$,

$$\begin{aligned}
g'(a(y))a'(y) &= -\frac{(\rho + \nu)V''(y)}{\beta\bar{y} \left[\left(1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V'(y)) - y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) V''(y) \right]} \\
&\quad + \frac{(\rho + \nu)V'(y)D}{\beta\bar{y} \left[\left(1 - (1 + \omega) \left(\frac{y}{\bar{y}}\right)^\omega\right) (\psi - V'(y)) - y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right) V''(y) \right]^2} \\
&= \underbrace{(\rho + \nu)V'(y)}_{>0} \frac{(V'(y) - \psi)\omega(1 + \omega)\frac{y^{\omega-1}}{\bar{y}^\omega}}{\{>0\}},
\end{aligned}$$

where $\{> 0\}$ denotes a strictly positive term and where we use the fact that $V''(y) = V'''(y) = 0$ for all $y \in (y^a, y^b)$. Since g is increasing the preceding equality implies that $a'(y) < 0$ for all $y \in (y^a, y^b)$.

Differentiating the right-hand side of the government's first-order condition with respect to y yields a strictly negative expression for all $y \in (y^a, y^b)$, because of condition (23), $g(a(y)) > 0$, and $V''(y) = 0$. Accordingly, the right-hand side of equation (24) is strictly decreasing in y , and so must be the left-hand side, which requires $a'(y) > 0$ for all $y \in (y^a, y^b)$. We have thus arrived at a contradiction and conclude that there exist no open intervals on the domain \mathcal{Y} such that $V''(y) = 0$ for all points in the interval.

[3] We conclude from [1] and [2] that V is strictly convex on the domain \mathcal{Y} . \square

G Observable Infection Status

When the infection status is public information all unaware households in the pre, post, and neutral groups avoid contact with the aware households in the post group. Since households in the pre group only interact with unaware members of the post group the law of motion reads

$$f(y, a) = g(a)\beta\bar{y}(1 - \sigma)y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right). \quad (25)$$

All unaware households have the same activity level a . This law of motion also applies under the Ramsey policy because it is optimal for the government to separate aware households from the rest of the population. In appendix G we analyze the government's program and the equilibrium.

When the infection status is private information⁵⁸ all post households interact with pre households. Pre households and unaware post households choose the activity level $a(y)$, aware post households choose a^* . The law of motion thus reads

$$f(y, \bar{a}) = g(\bar{a})\beta\bar{y}y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right), \quad (26)$$

where $\bar{a} \equiv (\sigma ya^* + (\bar{y} - \sigma y)a)/\bar{y}$ denotes average activity.

With public information the government's program is represented by the HJB equation

$$(\rho + \nu)V(y) = \max_a \sigma y u(a^*) + (1 - \sigma y)u(a) - \psi f(y, a) + f(y, a)V'(y) + \nu U^*$$

subject to the law of motion (25). Compared with the baseline case the benefit of activity is scaled by $1 - \sigma y$ because the activity level of aware households is exogenous, and the cost of activity is multiplied by the factor $1 - \sigma$ because of the modified law of motion. Accordingly, the first-order condition reads

$$(1 - \sigma y)u'(a(y)) = (\psi - V'(y))g'(a(y))\beta\bar{y}(1 - \sigma)y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right),$$

⁵⁸COVID-19 symptoms often are private information: Loss of smell and taste is unobserved by third parties and coughing is uninformative about the respiratory disease causing it.

which differs from the condition in the baseline case only insofar as β is replaced by $\beta(1 - \sigma)/(1 - \sigma y) < \beta$. Under the functional form assumption 4 the optimal activity level for unaware households therefore satisfies condition (6) except that β is replaced by $\beta(1 - \sigma)/(1 - \sigma y)$.⁵⁹

Turning to laissez faire the environment of unaware households with value function $U(y)$ differs twofold from the situation in the baseline scenario. On the one hand aggregate infection flows solely affect the unaware pre group; members of that group face the infection flow $f(y, a)/(1 - \sigma y)$. On the other hand infections do not only generate the deterministic costs considered so far but they also stochastically release affected households with flow probability σ into the pool of aware households generating a private capital gain $U^* - U(y)$.

Formally, under assumption 3 unaware households solve

$$\max_{a_i} u(a_i) + \frac{a_i}{a(y)} \frac{f(y, a(y))}{1 - \sigma y} (-\zeta\psi + \sigma(U^* - U(y)))$$

subject to the law of motion (25) where $a(y)$ denotes the equilibrium activity choice of all unaware households. The first-order condition is given by

$$(1 - \sigma y)u'(a_i) = \frac{\zeta\psi - \sigma(U^* - U(y))}{a(y)} g(a(y))\beta\bar{y}(1 - \sigma)y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right).$$

Under our functional form assumption 4 this implies that $a(y)$ satisfies condition (9) except that $\zeta\psi$ is replaced by $\zeta\psi - \sigma(U^* - U(y))$ (reflecting the capital gains reaped by aware households) and β is replaced by $\beta(1 - \sigma)/(1 - \sigma y)$. Compared with the optimal activity choice we see that the differences that were present in the baseline model (and that were reflected in static and dynamic externalities) continue to be present. But in addition the capital gains of households that become aware of their infection status now introduce a third difference. In equilibrium, these capital gains introduce an incentive for unaware households to engage more strongly in activity.

With private information, the government's program is identical to the program in the baseline model. This follows from the fact that the government cannot gain from letting aware households choose a different activity level than unaware types, even if this were feasible given households' incentive compatibility constraints.⁶⁰

In equilibrium, in contrast, aware and unaware households optimally choose different activity levels. Upon learning about their infection aware households attain utility U^* , as with public information.⁶¹ The value function of an unaware household in the pre or post

⁵⁹Average activity is given by $\sigma y a^* + (1 - \sigma y)a(y)$. In equilibrium, discussed below, it is given by the same type of weighted average.

⁶⁰To see this, suppose that unaware and aware households had activity levels a_1 and $a_2 > a_1$, respectively, such that $\bar{a} = (\sigma y a_2 + (\bar{y} - \sigma y)a_1)/\bar{y}$. The direct utility gain relative to the situation with a common activity level a_1 would equal $\sigma y(u(a_2) - u(a_1)) \approx u'(a_1)(a_2 - a_1)\sigma y$, while the cost would be equal to $(\psi - V'(y))(g(\bar{a}) - g(a_1))\beta\bar{y}y(1 - (y/\bar{y})^\omega) \approx (\psi - V'(y))g'(a_1)(a_2 - a_1)\beta\bar{y}y(1 - (y/\bar{y})^\omega)\sigma y/\bar{y}$. The first-order condition for a common activity level implies $u'(a_1) = (\psi - V'(y))g'(a_1)\beta\bar{y}y(1 - (y/\bar{y})^\omega)$. Since $\bar{y} < 1$ it follows that the losses from a discriminatory policy would exceed the gains.

⁶¹This is consistent with the private information assumption when insurers must not disclose their clients, perhaps for contractual or legal reasons.

group, $U(y)$, mirrors the function in the public information case except that unaware households perceive the modified law of motion. Under assumption 3 the first-order condition thus reads

$$(1 - \sigma y)u'(a_i) = \frac{\zeta\psi - \sigma(U^* - U(y))}{a(y)}g(\bar{a}(y))\beta\bar{y}y \left(1 - \left(\frac{y}{\bar{y}}\right)^\omega\right).$$

Imposing functional form assumption 4 and the equilibrium requirement $a_i = a(y)$ this yields

$$a(y) = \begin{cases} \frac{1 - \Sigma(y)\Theta(y)}{1 - \Sigma(y)\Theta(y) + \Theta(y)} & \text{if } n = 1 \\ \frac{-1 - 2\Theta(y)\Sigma(y)(1 - \Sigma(y)) + \sqrt{1 + 4\Theta(y)(1 - \Sigma(y))}}{2\Theta(y)(1 - \Sigma(y))^2} & \text{if } n = 2 \end{cases},$$

where we let $\Theta(y) \equiv (\zeta\psi - \sigma(U^* - U(y)))\beta\bar{y}y(1 - (y/\bar{y})^\omega)/(1 - \sigma y)$ and $\Sigma(y) \equiv \sigma y/\bar{y}$.