Journal of Economic Dynamics & Control 178 (2025) 105145

Contents lists available at ScienceDirect

Journal of Economic Dynamics and Control

journal homepage: www.elsevier.com/locate/jedc

A tractable model of epidemic control and equilibrium dynamics ☆

Martín Gonzalez-Eiras [a], Dirk Niepelt [b,*]

[a] University of Bologna, Piazza Scaravilli 2, 40126 Bologna, Italy
[b] University of Bern; CEPR, Schanzeneckstrasse 1, 3012 Bern, Switzerland

ARTICLE INFO

JEL classification:
D62
I18

Keywords:
Epidemic
Lockdown
Stimulus
Logistic model
Optimal control
COVID-19

ABSTRACT

We develop a single-state model of epidemic control and equilibrium dynamics, and we show that its simplicity comes at very low cost during the early phase of an epidemic. Novel analytical results concern the continuity of the policy function; the reversal from lockdown to stimulus policies; and the relaxation of optimal lockdowns when testing is feasible. The model's enhanced computational efficiency over SIR-based frameworks allows for the quantitative assessment of various new scenarios and specifications. Calibrated to reflect the COVID-19 pandemic, the model predicts an optimal initial activity reduction of 38 percent, with subsequent stimulus measures accounting for one-third of the welfare gains from optimal government intervention. The threat of recurrent infection waves makes the optimal lockdown more stringent, while a linear or near-linear activity-infection nexus, or strong consumption smoothing needs, reduce its stringency.

1. Introduction

Sudden epidemiological shocks such as the onset of the COVID-19 pandemic expose societies to severe tradeoffs between economic activity and public health. Early on in an epidemic, these tradeoffs are highly uncertain as it takes time to learn about epidemiological characteristics, the arrival rate of an effective vaccine or cure, and other factors determining likely scenarios. This uncertainty complicates a robust and effective policy response. What is needed is a flexible and computationally efficient framework to easily evaluate policy responses in a variety of scenarios.

No such framework is currently available. The large 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, authors routinely adopt strong (but different) assumptions they deem necessary to allow for numerical simulations.[1] As a consequence, the literature contains a large set of distinct, mostly quantitative findings, and commonalities are rare and obscured.

We propose an alternative framework with a single epidemiological state variable. Leveraging its simplicity and the resulting computational efficiency, we identify key mechanisms and characterize optimal policy across a wide range of scenarios. While the one-state

☆ For comments and discussions we thank the editor, Thomas Lubik, two anonymous referees, Alexander Ludwig and seminar participants at CEMFI, Universidad Torcuato Di Tella, Universities of Birmingham and Copenhagen, ANCE, DTMC COVID-19 Workshop, SED meeting, VfS Macro meeting, and Virtual Macroeconomics Seminar Series.

* Corresponding author.
E-mail addresses: mge@alum.mit.edu (M. Gonzalez-Eiras), dirk.niepelt@unibe.ch (D. Niepelt).
URLs: https://alum.mit.edu/www/mge (M. Gonzalez-Eiras), https://www.niepelt.ch (D. Niepelt).

[1] Models with two endogenous state variables typically encounter numerical problems and require restrictive assumptions to deal with them. For example, Alvarez et al. (2021) have to assume a counterfactually high share of the initially infected population. Farboodi et al. (2021) report "widely divergent paths" when solving backwards from the endogenous terminal state based on a chain of piecewise diverging subpaths; they need to impose a specific elasticity for the infections-activity nexus.

https://doi.org/10.1016/j.jedc.2025.105145
Received 3 October 2024; Received in revised form 11 July 2025; Accepted 11 July 2025

Available online 17 July 2025

simplification has drawbacks—most notably, it implies that policy and long-run health outcomes are orthogonal—we demonstrate that this is irrelevant for the optimal policy during the critical early stages of an epidemic, when SIR-based frameworks also exhibit near orthogonality. Indeed, under a baseline parameterization and scenario, the optimal policy paths in our model and a standard SIR-based framework virtually *coincide* for many months. Accordingly, we view our model as a robust and flexible tool for analyzing trade-offs at the onset of an epidemic.

The generic epidemiological framework we use generalizes the "simple epidemiological model" (Bailey, 1975) and captures the essence of infection 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—our single endogenous state variable—approximately follow a logistic law of motion, which accurately approximates dynamics in two-state SIR models as we show.

The economic layer that we superimpose on this epidemiological structure incorporates households and a government. As is standard in the literature, 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 infections, which are privately and socially costly. Households are fully aware of aggregate infection dynamics, behave individually rationally, and in equilibrium shoulder the entire social costs of infection. Nevertheless, they fail to fully internalize the consequences of their activity choices, and this gives rise to "static" and "dynamic" externalities.

Such externalities call for government intervention (Gersovitz and Hammer, 2004). We find that this goes both ways. We prove that the optimal intervention features lockdowns in the beginning of the epidemic followed by "inverse lockdowns," namely interventions to stimulate private sector activity. Intuitively, as the epidemic progresses, there comes a point at which individual households act too *cautiously*, thereby unwillingly delaying the end of the epidemic. At this point, the government optimally stops curbing activity and starts promoting it. Our result may rationalize measures imposed during the COVID-19 pandemic such as monetary easing, temporary sales tax reductions, subsidies, or "return-to-work bonuses."[2]

Next, we allow lockdowns to be accompanied by testing policies, which identify infectious individuals with a certain probability. When the government has the power to temporarily quarantine such individuals, testing slows down infections and is valuable, rendering lockdowns and testing substitutes. Conversely, the prospect of future testing capabilities renders lockdowns and testing complements: The government optimally curtails activity to buy time for these capabilities to become available.

We show that the situation is quite different under laissez faire. If privately administered tests make individuals aware of the fact that they gained immunity, then they make these individuals return to work full time, raising the infection risk for others, and motivating all those unaware of their infection status to turn more and more cautious as the epidemic progresses. What is more, the prospect of learning about one's immunity status, and thus returning to work full time, introduces an upside in the program of households and reduces the incentive to take precautions. In fact, we prove that in the early stages of an epidemic, testing fosters equilibrium activity.

Our third set of findings comprises quantitative results. Exploiting the fact that the model can easily be numerically solved and simulated, we derive the optimal policy under general assumptions about the elasticity of the activity-infection nexus—unlike much of the literature, which typically assumes an elasticity of one or two to avoid problems with the numerical solution. This flexibility allows us to address the mixed epidemiological evidence on the elasticity (Hethcote, 1989). It also opens the possibility to adapt the model's policy recommendations to regional (e.g., rural vs. urban) or other disparities.

We also study the role of other key parameter (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. Moreover, we analyze optimal policy under recurrent infection waves or when the government's instruments are limited such that inverse lockdowns become infeasible.

Calibrated to match COVID-19 infection data in the U.S., the baseline model indicates that from mid-March 2020, economic activity should have been immediately reduced by nearly forty percent. This policy would have increased welfare by approximately 0.32 percent of lifetime consumption (corresponding to around 2500 U.S. dollars per capita in present value) compared to a laissez-faire approach. Across all the specifications and scenarios we examine, many yield expected results. The optimal lockdown lengthens (in state space), and activity falls, when the fatality rate or the arrival rate of a cure increase, when congestion effects are stronger, or when learning effects are weaker. Testing policies that reduce the effective infection rate, or the prospect of such policies, increase the lockdown duration as well.

Other simulation results, specifically some quantitative effects, are more surprising. The threat of recurrent infection waves triggers a large increase in the optimal lockdown duration and strictness, while a linear or near-linear activity-infection nexus or strong consumption smoothing needs have the opposite effect. The model implied welfare gains suggest that government intervention is particularly important when recurrent infection waves are a threat. The most striking simulation result concerns inverse lockdowns—stimulus measures the government imposes once the economy starts to see the end of the tunnel: It is these measures that are responsible for a large, maybe even the dominant share of the welfare gains from optimal government intervention.

Crucially, our model replicates the predictions of a prototypical model based on the two-state SIR framework when we calibrate it correspondingly. Fig. 1 plots the optimal activity level (which varies between zero and one) in Farboodi et al. (2021) and in our model against the cumulative share of the infected population (y).[3] During the initial phase of the epidemic, the two predictions essentially

[2] In the U.S. the National Economic Council urged lawmakers to replace a lump-sum transfer with a "return-to-work bonus" (*The Wall Street Journal*, June 15, 2020). In the U.K. during August 2020 the government implemented a program aimed at encouraging people to eat in restaurants (Fetzer, 2022).

[3] See appendix A for a detailed discussion.

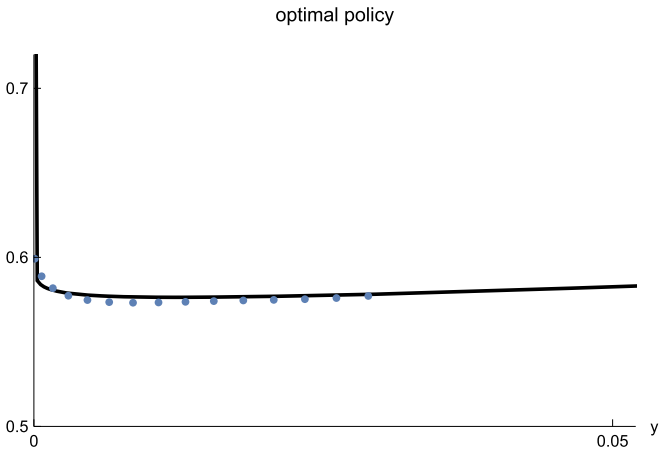

Fig. 1. Optimal policy: Logistic model (solid) vs. two-state SIR model (dotted).

coincide. Intuitively, both the SIR framework's added complexity of epidemiological dynamics and the resulting endogeneity of long-run health outcomes play practically no role for optimal policy in the early stages of an epidemic. The simplicity and computational efficiency of our model therefore come at minimal cost during these early stages.

Our final contribution is of a technical nature. We establish that the value function is differentiable over the entire state space, implying that optimal policy is continuous and standard numerical methods may justifiably be applied to simulate the model. This puts our framework with one endogenous state variable in stark contrast to much of the macro-epidemiological literature, in which the derivation of analytical results is challenging and the applicability of numerical methods questionable because the law of motion of infections renders the Hamiltonian function of the control problem non-convex (see, e.g. Calvia et al., 2023).

Related literature Workhorse epidemiological models are due to Kermack and McKendrick (1927) and Bailey (1975); for reviews, see, e.g., Hethcote (1989) and Hethcote (2000). The COVID-19 pandemic has spurred broad interest in economic cost-benefit analysis in the context of epidemiological dynamics. Early contributions include Atkeson (2020) and Eichenbaum et al. (2021). Alvarez et al. (2021) compute the optimal lockdown policy to flatten the infection curve in order to relax health care capacity constraints. Farboodi et al. (2021) argue that in equilibrium and under the optimal policy the effective reproduction number remains close to unity.

Gersovitz and Hammer (2004), Bethune and Korinek (2020), and Jones et al. (2021) assess externalities of privately optimal precautions in the epidemiological context. In SIR based parallel work Garibaldi et al. (2024) decompose these externalities into static and dynamic parts; see also Rachel (2022). We show that dynamic externalities necessarily start negative and eventually turn positive; that static externalities are present even when the infection matching function exhibits constant returns to scale; and we derive the implications of these facts for optimal "inverse lockdowns."

Piguillem and Shi (2022) simulate a scenario with widespread testing. They find that even non-targeted tests are an effective and cost-efficient substitute for lockdowns. Pollinger (2023) assumes that infections stop once the infectious pool falls below a critical mass and characterizes optimal suppression policies with or without targeted tests. He finds that testing allows to relax lockdowns. We analytically show that testing policies may substitute for lockdowns even in the absence of such a critical mass; and we contrast this substitution effect with complementarities that arise from the prospect of future testing capabilities.

Kaplan et al. (2020), Acemoglu et al. (2021), and Ellison (2024) analyze the implications of heterogeneity, including differential costs of reduced activity, asymmetric infection dynamics due 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. We focus on uniform lockdown policies to study the government's program in a variety of scenarios and extensions. Moreover, we show that in spite of the single endogenous state variable, our analysis allows to capture some degree of heterogeneity within groups.

Most of this work focuses almost exclusively on numerical analyses, with Abel and Panageas (2020), Gonzalez-Eiras and Niepelt (2020a), Miclo et al. (2022), Rachel (2022), and Toxvaerd (2020) constituting some notable exceptions.[4] Our paper combines novel analytical results with numerical simulations.

Following Richards (1959) many empirical studies apply the generalized logistic model to epidemic dynamics. For example, Wu et al. (2020) simulate COVID-19 infection dynamics in several regions based on the generalized logistic model and they note the short-term forecast accuracy of the model. They abstract from behavioral responses, unlike our model.[5]

[4] 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. Rachel (2022) argues that externalities can be negative and optimal policy may avoid recurrent infection waves.

[5] Wu et al. (2020) find that "second waves" are important in the data but not captured by their model. Our numerical analyses allow for recurrent waves.

Outline The remainder of the paper is organized as follows. We lay out the model in section 2 and present the conditions characterizing first best and equilibrium in section 3. In sections 4 and 5, we characterize externalities, lockdowns, inverse lockdowns as well as the implications of testing policies analytically. The quantitative analysis of a series of model modifications and extensions is contained in section 6. Section 7 concludes. Auxiliary discussions are relegated to appendices; proofs of lemmas and propositions in the main text are 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.[6] Our framework incorporates one endogenous state variable (rather than two in the typical SIR model) and possibly time as a second state variable (unlike SIR and SIS models). Below, we will also introduce 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" 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,[7]

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

Accordingly, the share of the pre group changes by $dx/dt = -dy/dt$. Variable a in equation (1) is 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.[8] 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 then[9]

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

The law of motion (1) generalizes the SI model (Bailey, 1975), incorporating the effect of economic activity on the infection rate.[10] It is also closely related to the SIS model, where infected individuals randomly recover and return to the susceptible pool (Hethcote, 1989).[11] Our framework differs from the SIS model in that we model infections as flows, $f(t)$, whereas the SIS model focuses on the stock of infected individuals.[12]

Most importantly, the law of motion (1) is closely related to both the canonical SIR model (Kermack and McKendrick, 1927) and the modified SIR model (Bailey, 1975). In their general form, both models describe the evolution of three population groups—susceptible, infected, and removed (recovered or deceased)—based on two endogenous state variables (the shares of two of these

[6] The "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.

[7] As is standard, we model epidemic dynamics in continuous time. This is not only consistent with the initial formulation of the SIR model (Kermack and McKendrick, 1927) and as a consequence, most macro-epidemiological frameworks in the literature, but it also avoids complex or even chaotic dynamics that can arise with a discrete logistic function (see, e.g., May, 1974).

[8] 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, introduced below, 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 total utility by $1 - \delta y \approx 1$.

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

[10] The SI model assumes $\omega = \bar{y} = 1$.

[11] The SIS model assumes $\bar{y} = 1$.

[12] In Gonzalez-Eiras and Niepelt (2023), we allow for members of the post group to return to the pre group, for instance because infection does not confer permanent immunity.

three groups). Susceptible households become infected through contact with currently infected households, remain infectious for a random duration, and eventually transition to the removed group.

These dynamics reduce to the law of motion (1) with $\omega = 1$, when the infected and removed groups are combined into a single post group, effectively assuming a fixed relative size of the infected and removed populations.[13] This key simplification eliminates an endogenous state variable but does not undermine the model's ability to represent societal costs of infection or death. For these purposes, it is sufficient to account for the *flow* of newly infected individuals from the pre to the post group, and to associate costs with this flow, as outlined below.[14]

The steady state of system (1), $\bar{y}$, is exogenous, whereas it is endogenous in the canonical (but not the modified) SIR model. However, this difference is less significant than might initially appear. Since we allow for a random arrival of a cure (see below), the expected long-run share of the population that avoids infection is endogenous to the chosen activity path, as in the SIR model. Furthermore, as shown in appendix A, the endogeneity of the "herd-immunity" threshold in the canonical SIR model does not significantly impact optimal policy choices during the early stages of an epidemic.[15] During this initial phase, our framework therefore constitutes a useful environment for policy analysis that is not yet technically feasible within the canonical SIR model.[16]

In addition to incorporating the effect of economic activity on the infection rate, our framework extends the SI, SIS, and SIR models by introducing time-dependent dynamics, in addition to the state-dependent ones. With Poisson arrival rate ν, a "cure arrives"—marking the point at which the disease and its associated consequences (described below) disappear, and β drops to zero. This cure could represent, for instance, medical advancements or the development of a vaccine. The same cure also arrives deterministically at a fixed time, T, which may be infinite.

2.1.2. Costs of infection

Infections are transitions of households from the pre to the post group. These transitions generate social costs, due to strain on the healthcare system, foregone utility, or harm caused by the infection. We express these costs as

$$\psi\, f(y,a), \tag{3}$$

where $\psi > 0$ denotes the unit cost. In the extensions discussed in section 6.2, we allow ψ to vary in order to capture effects such as learning or congestion. In the former case, unit costs decrease as cumulative case numbers increase. In the latter case, unit costs rise with the infection flow, for example, due to healthcare capacity constraints that lead to lower care quality and higher fatality rates—creating an incentive to "flatten the curve." In the baseline analysis, we assume ψ is constant.

We emphasize the similarities between the drivers of health costs in SIR based analyses and in our framework: In the former, an exogenous share of the infected population dies at each instant, and public health costs are associated with these transitions into death. Due to the exogenous death rate, this is equivalent to associating costs with transitions into the infected pool, which evolves endogenously. In our model, the public health costs also are associated with the transition into the post pool, and this pool also evolves endogenously.

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 $f : [0, \bar{y}] \times A \to \mathbb{R}^+$ given in equation (1) determines epidemiological dynamics. Function $g : A \to \mathbb{R}^+$ 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.

2.2. Economics

2.2.1. Households

Households have an intertemporal utility function that depends on activity and infection costs. In combination with the law of motion (1), these preferences parsimoniously generate the fundamental tradeoff of interest between lives and livelihoods. Households discount the future at rate ρ.

[13] In both the SIR model and our framework, infections are caused by interactions between the susceptible (pre) group and the infectious group. In our model, the latter group constitutes a fixed share of the post group. The coefficient β absorbs the ratio between infectious and post groups.

[14] The law of motion (1) does not explicitly account for deaths. However, the implications for epidemiological dynamics are negligible when death rates are assumed to be small, as we do here. See the discussion in footnote 8.

[15] A hybrid model augments the modified SIR model with an additional parameter that regulates long-run population shares (Gonzalez-Eiras and Niepelt, 2020c). In the law of motion (1), $\bar{y}$ serves a similar role: a lower $\bar{y}$ implies a larger proportion of the population that remains uninfected.

[16] We focus on the model's predictions in y space rather than in the time domain. (Of course, there is a direct mapping between the two.) The y-space perspective provides more insight because y rather than time is the key state variable in the model. For our COVID-19 calibration, with a basic reproduction number of $\mathcal{R}_0 = 2.4$, a conservative lower bound for the initial phase corresponds to a cumulative infection rate y of 3.5 percent of the population. Simulation results show that for higher values of $\mathcal{R}_0$, the initial phase extends over a larger portion of the state space.

The benefit of activity is represented by an indirect utility function, u, which depends on the level of activity and reaches a maximum at the first-best activity level, $a^\star$. Without loss of generality, we normalize u and g such that $a^\star = 1$, $u(a^\star) = 0$, and $g(a^\star) = 1$.

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

The cost of activity is increased risk of infection, scaled by infection costs. To ensure internal consistency, we require that the perceived relationship between individual activity choice and infection risk aligns with the epidemiological environment in which individuals operate—namely, the aggregate law of motion for infections (1)—so that equilibrium behavior reflects rational, forward-looking decision-making. In section 3.2, we describe the mapping between individual activity choice and infection risk and the resulting optimization problem of an individual household in decentralized equilibrium. We analyze heterogeneity in section 5.

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 section 6.6 we analyze the consequences of instrument restrictions, which imply that the government faces a more limited choice set than a social planner.

The government treats households symmetrically. This could reflect that the government has no information about individual health status, for example because symptoms do not generate much relevant information about their cause. It could also reflect that the government does have such information but chooses to disregard it because conditioning policy on health status would be too costly (for reasons the model does not speak to). Note that the government can condition policy on the aggregate state y even if it does not observe the health status of each individual, at a minimum because knowing y_0 and the law of motion (1) allows to infer the aggregate state.

Since the losses from infection are linear in both infection risk and the health costs conditional on contracting the disease, they can be summed as total losses. Accordingly, the government's objective incorporates the summed losses rather than a nonlinear aggregate of individual losses. Formally, the government's program introduced in subsection 3.1 treats all households as bearing an equal share of the social costs—regardless of their infection status—which equals one in equilibrium. Moreover, all household types—pre, post, and neutral (the latter being unaware of their status)—are prescribed the same activity level.

The symmetry assumption is a plausible approximation in the context of many epidemics, including the COVID-19 pandemic.[17] It is clearly 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 section 5 the implications of heterogeneity.

2.3. Functional form assumptions and calibration

To sharpen our results we sometimes impose the preference assumption

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

consistent with the normalizations $a^\star = 1$ and $u(a^\star) = 0$. Our preferred interpretation is that activity yields strictly concave benefits (e.g., from consumption) and linear costs (e.g., from providing effort). In section 6.7 we 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 3. 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.[18] 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).[19]

[17] Many 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.

[18] Mathematica code is available upon request.

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

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

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 (introduced in section 3.2)

In appendix C we describe in detail how we calibrate the remaining parameters. Based on information about parameter values in the canonical SIR model we let $y_0 = 0.1148 \cdot 10^{-3}$, $\beta = 0.8346 \cdot 10^{-1}$ (corresponding to an infection rate in the SIR model (at regular activity level) of 0.1333), $\omega = 0.6662$, and $\bar{y} = 0.8786$. Table 1 summarizes the baseline calibration.

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 $U^\star \equiv u(a^\star)/\rho$ 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 $U^\star$, at time $t < T$, the state in the programs of the government or a household is given by (y, d) where $d \equiv T - t$ indicates duration until T.[20] 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 change in the infinite-T case.

3.1. Government program

The government weighs the social costs of reduced economic activity against the benefits of delaying infections and potentially avoiding major suffering from the disease before a cure becomes available. As explained earlier, our parsimonious specification represents the cost of activity reductions by a lower value of $u(a)$ and health cost of infections by the term $\psi f(y, a)$. The government knows the aggregate state and the law of motion (1) and uses this information to optimally trade off lives and livelihoods. Since this tradeoff is dynamic, the optimal policy solves an optimal control problem. While this control problem is straightforward in principle we devote some attention to seemingly "technical" aspects such as differentiability of the value function and continuity of the optimal policy function, because they play an important role for central properties of the optimal policy, as we show.

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

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

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), $P^\star$, depends on parameters including T. It represents the expected present value of $U^\star$ realizations either at time T, or after Poisson shocks that occur before time T.[21] 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).

[20] Following standard practice, we formulate the law of motion (1) using calendar time, t, rather than duration, d. In contrast, we find it more intuitive to use d rather than t as an argument of the value function. Of course, the two alternatives are intimately connected.

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

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

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. Let $D_y V$ 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.[22]

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

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

with boundary condition $V(y,0) = U^\star$; moreover, $V(y,d) < U^\star$ 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 $V(0) = V(\bar{y}) = U^\star$; moreover, $V(y) < U^\star$ for all $y \in (0,\bar{y})$ and V is Hölder continuous with exponent $\min[\frac{\rho+\nu}{\beta \bar{y} \psi \max[\omega,1]}, 1]$.

In appendix D we briefly review the concept of viscosity solutions for non-linear partial differential equations (Crandall and Lions, 1983). Interpreting $V(y,d)$ as the value of an asset that the government optimally manages, the left-hand side of the HJB equation represents the "required return" on that asset, due to time discounting. The right-hand side captures the two sources of return, the "dividend" and "capital gains" components. The dividend component stems from the immediate utility flow, which equals the indirect utility from activity, net of infection costs, $u(a) - \psi f(y,a)$. The capital gains come from the evolution of the state or the sudden arrival of a cure. A change of y (given by $f(y,a)$) induces the change of value $D_y V(y,d)$, a reduction in d the gain $-V_d(y,d)$, and the arrival of the cure triggers the gain $U^\star - V(y,d)$.[23]

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 D). The corresponding optimal control, $a(y,d)$, then solves

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

or

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

Equation (5) represents the two opposing forces that shape the government's choice of activity level: economic needs and public health concerns. The direct effect of economic costs appears on the left-hand side, while public health costs are reflected in the first term on the right-hand side. The second term on the right-hand side captures the change in the present value of future economic and public health losses. Note that faster infections cause higher health costs, captured by the ψ term, and change the endogenous state, bringing the epidemic closer to its end.

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 3, condition (5) reduces to the policy function[24]

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

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 establishes 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*

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

[22] Recall that a 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 $x_1, x_2 \in X$. 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 $x_1, x_2 \in X$. When $k < 1$ Hölder continuity is a strictly weaker requirement.

[23] See, for example, Acemoglu (2009, ch. 7) for the interpretation of the HJB equation in terms of dividends and capital gains.

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

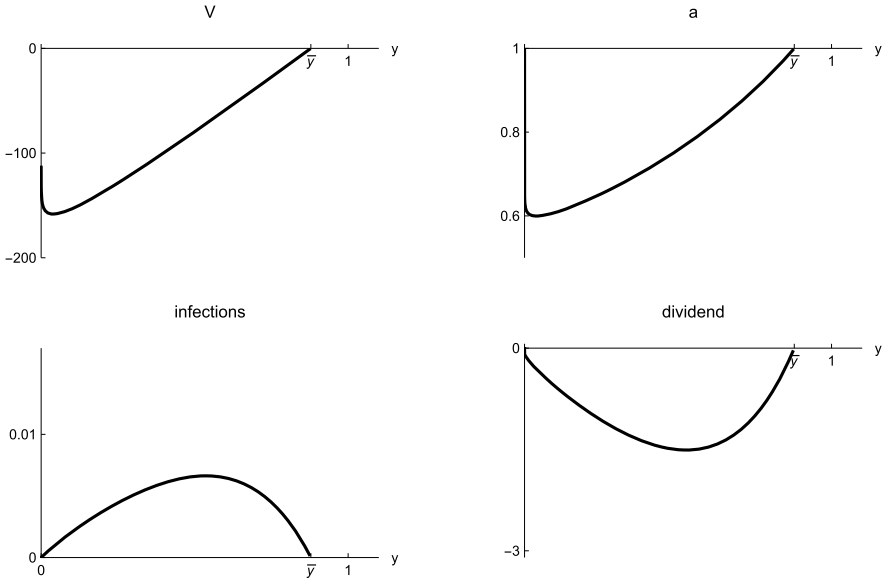

Fig. 2. Value function, activity level, infections, and dividend in the government's program.

Moreover, $\lim_{y\downarrow 0} a(y,d) = a^\star$. *When* $T = \infty$, *these results hold for* $V(y)$ *and* $a(y)$. *A change of any parameter other than* ρ, ν, *or* T *that increases* $V(y,d)$ *also increases* $a(y,d)$.

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 3, the proposition implies $(\rho+\nu)V(y,d) = \ln(a(y,d)) + (1-a(y,d))(1-n^{-1})$; 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 $V(0) = V(\bar{y}) = U^\star$ and $V(y) < U^\star$ 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 E we show that V has a unique minimum, 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 boundary condition $V(\bar{y}) = U^\star$ and use finite difference methods in Mathematica, see appendix C.3. We use a parallel strategy to characterize the decentralized equilibrium discussed below.

Fig. 2 illustrates our results.[25] The figure is drawn for the baseline calibration introduced earlier and under the assumption that $g(a) = a^2$. In the quantitative analysis in section 6, we will conduct various robustness checks and consider alternative specifications and scenarios.

The results in Fig. 2 are intuitive. Focus first on the optimal path of infections displayed in the bottom left panel. Infections follow a hump-shaped pattern—consistent with standard epidemiological predictions—because the generalized logistic function governing the evolution of y is S-shaped, making infection flows (i.e., changes in y) largest at intermediate values of y. The endogenous variation in $a(y)$ controlled by the government does not offset this fundamental force. 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: Higher infections are associated with more negative dividends.[26]

The capital gains component, $f(y,a(y))V'(y) + \nu(U^\star - V(y))$, is the difference between the required return, $\rho V(y)$, and the dividend component. Since the slope of V (displayed in the top left panel) is rather flat for $y > y^{\min}$, the first part of the capital gains component is positive and hump-shaped in this domain, while the second part is positive and decreasing. The first part reflects the economy's nearing the end of the tunnel: Infections today means fewer infections in the future, which is beneficial because fewer future infections are associated with a reduction in the present value of future health and inactivity costs. The second part of the capital gains component captures the benefit of the potential arrival of a cure, which is diminishing as the epidemic progresses.

For $y \le y^{\min}$, the first part of the capital gains component is negative: While the economy moves toward the end of the tunnel, it still has to navigate through its darkest part. This explains why the value function is falling so steeply. The second part of the capital gains (which is not under the government's control and thus will not directly shape optimal policy) increases because the benefit of the potential arrival of a cure reaches its maximum when the value of V is minimal.

[25] It is understood that these relationships apply before a cure arrives. Once a cure has arrived a optimally equals $a^\star$ even if $y < \bar{y}$.

[26] Dividends are negative because $u(a^\star)$ is normalized to zero.

The optimal activity path displayed in the top right panel reflects the tradeoff between lives and livelihoods described after equation (5). The state y affects optimal activity through the difference between infection costs and the first part of the capital gains component, $f(y,a(y))(\psi - V'(y))$.[27] Via its effect on f, the S-shaped infection path strengthens the motive to reduce a for intermediate values of y, while the term $\psi - V'(y)$ is very large for small values of y before flattening out. In combination, the two effects lead the government to curtail activity strongly in the early phase of the epidemic and increasingly less so later on.[28] Formally, the relation between V and a established in Proposition 1 implies that optimal activity bottoms out at the same time as V, at $y^{\min}$, and exhibits very similar skewness.

It is instructive to consider two extreme cases in this context. When $\rho + \nu$ converges to zero such that the government does not discount the future, the derivative of the value function converges to ψ for all $y > 0$. As a consequence, the optimal activity level approaches $a^\star$ for all $y > 0$ (see equation (5)). Intuitively, when there is no hope for a cure and delaying infections does not help, then there is no point in trying to manage the time profile of infections, and thus in reducing activity.[29]

In contrast, when $\rho + \nu$ increases such that the government attaches less weight to the future or considers the arrival of a cure more likely, the value function flattens and approaches zero. In the limit, there are no capital gains, only current infection costs shape public health concerns, and the government suppresses activity more strongly than without discounting.

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 $a(\bar{y}) = 1$. We exploit this possibility in subsection 6.3.

3.2. Decentralized equilibrium

Unlike the government, an individual household takes the evolution of the aggregate state y as given. It optimizes subject to this aggregate constraint as well as the perceived nexus between the *individual* activity level and infection risk. To discipline our analysis, we impose that this perceived nexus is consistent with the epidemiological environment the individual actually inhabits, and thus with the economy-wide law of motion (1), allowing us to interpret equilibrium choices as rational, forward looking behavior.

Specifically, we assume that the infection cost function perceived by an individual household differs twofold from the social cost function faced by the government. On the one hand, the individual distinguishes between its own and the aggregate activity choice. Following standard practice (e.g., Farboodi et al. (2021) or Garibaldi et al. (2024)), we use a matching function to describe how this activity pair maps into private infection risk. On the other hand, a household may internalize only a share of the infection costs, for instance because of health insurance and non-contractable precautionary behavior.

We postulate a homogeneous of degree $n \geq 1$ matching function, m, that maps the household's individual activity level a_i as well as aggregate activity a into infections conditional on y. Consistency requires that $g(a) = m(a,a)$. A common matching function that satisfies our assumptions (and that we will use) is given by $m(a_i,a) = a_i^x a^{n-x}$ with $x > 0$; when $x = 1$ this yields $m(a_i,a) = (a_i/a)a^n$ such that an individual household perceives a linear effect of its activity choice.[30] Furthermore, we assume that a household internalizes the marginal share $\zeta \in [0,1]$ of infection costs but ends up bearing the full social costs in equilibrium. Formally, we impose the following assumption[31]:

Assumption 4. Households perceive the individual health cost function to be given by $\psi f(y,a)/g(a)$ multiplied by the factor $\zeta m(a_i,a) + (1-\zeta)g(a)$, $\zeta \in [0,1]$, where m is homogenous of degree n and $m(a,a) = g(a)$.

When the internalization rate ζ equals zero, households perceive the same cost function as the government (recall (3)), but since they distinguish between their own activity choice and the aggregate one, households take no precautions. When the internalization rate exceeds zero, in contrast, households perceive the marginal effect $\zeta(\partial m(a_i,a)/\partial a_i)\psi f(y,a)/g(a)$ of their activity on their own health costs; nevertheless, they bear the full costs in equilibrium, since $(\zeta m(a,a) + (1-\zeta)g(a))\psi f(y,a)/g(a) = \psi f(y,a)$.

We emphasize that this specification is not only consistent with common matching function specifications, as argued above, but also with rational, forward looking behavior. To see this, note that an individual that only knows the aggregate state considers herself to be a member of the pre group—and thus, at risk of infection—with probability $\bar{y} - y$. Conditional on being in the pre group, in turn, the risk of meeting an infectious individual is y. Finally, conditional on an infectious "match," the risk of actually contracting the disease is a function of the activity levels of the two parties, $m(a_i,a)$.[32] In what follows, we augment the functional form Assumption 3 to include the specification $m(a_i,a) = \frac{a_i}{a}a^n$, i.e., we set $x = 1$.

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 $\zeta = 0.8266$ (see appendix C).[33]

[27] The second part of the capital gains component does not interact with activity.

[28] The initial fall in optimal activity is barely visible in Fig. 2. But recall from Proposition 1 that $\lim_{y \downarrow 0} a(y) = a^\star$ and policy is continuous.

[29] Recall that $\bar{y}$ is exogenous.

[30] This formulation is consistent, e.g., with assumptions in Farboodi et al. (2021) or Garibaldi et al. (2024).

[31] This can be relaxed, see Gonzalez-Eiras and Niepelt (2020b).

[32] If the individual is not in the pre group, it faces no risk. The generalized logistic specification replaces the y in the $\bar{y} - y$ term with a geometric average.

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

Since an individual household (correctly) perceives the aggregate state to develop independently of its own choices, the household solves a static problem, trading off net economic activity benefits and expected infection costs. The household's choice solves

$$a_i(y,d) = \arg\max_{\alpha\in A} u(\alpha) - \{\zeta m(\alpha, a(y,d)) + (1-\zeta) g(a(y,d))\}\ \psi\ \frac{f(y,a(y,d))}{g(a(y,d))},$$

where $a(y,d)$ denotes the aggregate activity choice, which the household takes as given. In equilibrium, individual and aggregate choices coincide and thus satisfy

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

Any fixed point $a(y,d) \in A$ that solves this equation constitutes an equilibrium choice. This choice has several noteworthy properties. First, unlike in the government's program, equilibrium activity is symmetric around $\bar{y}/2$ when $\omega = 1$. This reflects the absence of capital gains in the household's program. Second, conditional on y, equilibrium activity is independent of time even when $T < \infty$, for similar reasons. 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 $(0,1)$.[34] For $n = 1, 2$ this solution is given by

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

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. Externalities, lockdowns, and inverse lockdowns

Next, we turn to the externalities that arise because households fail to fully internalize the societal consequences of their actions. Subtracting the right-hand side of the planner's first-order condition (5) from the right-hand side of the individual's optimality condition (8) and evaluating terms at a common activity level yields $\mathcal{S}(a,y) + \mathcal{D}(a,y,t)$ with

$$\mathcal{S}(a,y) \equiv \psi f(y,a)\left(\frac{\zeta}{a} - \frac{g'(a)}{g(a)}\right),$$

$$\mathcal{D}(a,y,t) \equiv f(y,a)\frac{g'(a)}{g(a)} V_y(y,t).$$

We refer to $\mathcal{S}(a,y)$ as the "static" externality and to $\mathcal{D}(a,y,t)$ as the "dynamic" externality.

The static externality arises because households do not fully internalize the contemporaneous costs of infection. It has two sources: First, if $\zeta < 1$, a household shifts the share $1-\zeta$ of the costs of its own infection onto society. Accordingly, it (correctly) perceives the private marginal cost of infection to be reduced by the factor $1-\zeta$. Second, if g is strictly convex, the household's perceived linear private activity-infection nexus contrasts with the non-linear, convex nexus at the aggregate level (recall the discussion preceding Assumption 4). Intuitively, aggregate infections exhibit increasing returns to scale with respect to (symmetric) activity in this case, which the individual household (correctly) does not perceive. Formally, with $g(a) = a^n$, the static externality is proportional to $(\zeta/n - 1)$ and thus strictly negative when $\zeta < 1$ or $n > 1$. We analyze the quantitative role of n in subsection 6.3.

The dynamic externality reflects the intertemporal wedge between private and social marginal rates of substitution. Since individual households (correctly) perceive their activity choice to have no effect on the aggregate state, which is a key statistic for continuation welfare, and because they behave accordingly, their choices fail to internalize the consequences for the continuation value. Note that both the static and the dynamic externality are proportional to infections, $f(y,a)$, because infection flows determine the immediate costs and change the state variable, which in turn affects the continuation value. Note also that the dynamic externality is positive, unlike the static one.

Fig. 3 illustrates the consequences of the externalities. As in Fig. 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 Fig. 2; the dashed lines represent the equilibrium outcomes.

Not surprisingly, the presence of externalities reduces the equilibrium value below the value under the government's optimal policy. For the same reason as in the government's problem (the worst part of the epidemic is still to come), the equilibrium value function falls steeply when y is very small. More interestingly, the activity levels displayed in the right panel differ markedly: Initially,

[34] By letting $\underline{a}$ be sufficiently small we can thus guarantee existence of equilibrium.

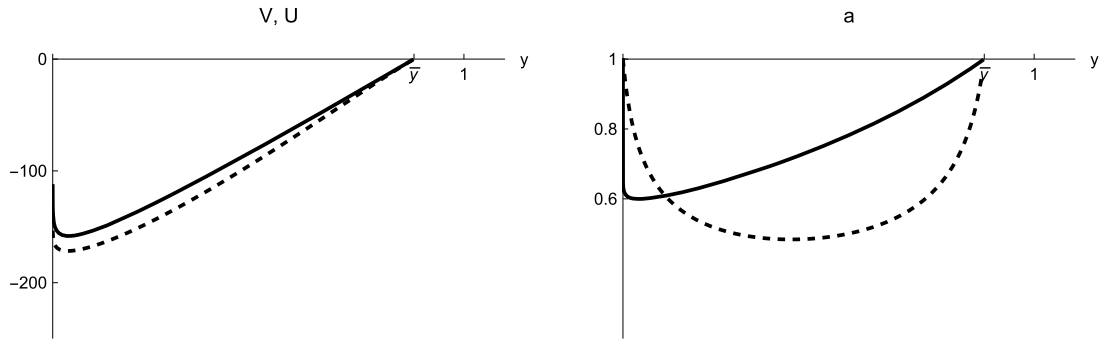

Fig. 3. Value function and activity level in the government's program (solid) and in equilibrium (dashed).

the government chooses a lower activity level than in equilibrium, but eventually the opposite holds true—households behave more prudently than socially optimal.

The driving force underlying households' excessive caution is the capital gains component (specifically its first, activity dependent part), 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 relative to the equilibrium outcome. Lockdown measures may include, for instance, stay-at-home-orders, social distancing rules, business closures, or school closures, while inverse lockdowns may, e.g., take the form of stimulus 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 $g(a)=a^n$).[35]

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 4 and if $g(a)=a^n$ and $\zeta \le n$, lockdowns occur as follows:*

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

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. It implies that, under a parametric condition, the dynamic externality eventually dominates the static one, leading the government to accelerate infections in order to reach the end of the tunnel more quickly.

Again, this implication is consistent with various stimulus measures during the COVID-19 pandemic, such as temporary tax reductions and employment or consumption subsidies. Our reversal result triggers the question how policy should optimally be conducted when the government cannot impose an inverse lockdown. We analyze this question in section 6.6, where we study the effect of constraints on the government's choice set.

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. At some point, society thus experiences capital gains, reflecting the change of an aggregate state variable (or multiple such variables). Since households take the evolution of the state as given, the capital gains are not internalized in equilibrium. As long as the capital gains are sufficiently large to outweigh negative static externalities, the reversal result thus follows.

Our results on static and dynamic externalities as well as lockdowns and inverse lockdowns relate to work by Rachel (2022) and Garibaldi et al. (2024). Rachel (2022) analyzes a setting with two possible activity levels and emphasizes the externality due to low activity choices by susceptible individuals. Garibaldi et al. (2024) employ a simplified SIR setting in discrete time.[36] Our results differ along several dimensions. We find that dynamic externalities are strictly positive towards the end of the epidemic, and we show that static externalities may be present even with $n=1$, because $\zeta<1$ also drives a wedge between private and social marginal

[35] Note 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.

[36] Since Garibaldi et al. (2024) consider two states (the masses of infected and susceptible households) they talk about two dynamic externalities. They also discuss congestion externalities, which our continuous time specification classifies as static.

health costs.[37] Moreover, we compare optimum and equilibrium activity levels in state space in order not to disguise fundamental differences in economic incentives.[38]

5. Testing and heterogeneity

A common complement to lockdowns are comprehensive test, trace, and quarantine policies with the objective to isolate infectious persons and break chains of infection. With slight modifications, our framework can speak to such policies.[39]

We assume that the intensity of testing is linked to new infections because it correlates with activity: An individual with higher activity is tested more often. There are several natural interpretations of this correlation, including that more intense testing requires more activity, or that more active individuals choose to test more often (or are required by government to do so), as is typically the case. Rather than modeling the resulting link between testing and infections, we impose it in reduced form.[40]

Suppose, then, that the test policy identifies new post individuals with conditional probability σ. Share σ of individuals in the post group thus know about their infection status—they are "aware"—and this status is common knowledge. Testing may also reduce the effective infection rate from β to $\tilde{\beta} = (1-\kappa\sigma)\beta$, where $0 \leq \kappa \leq 1$, for instance because of quarantine measures in connection with the tests.[41] In this environment, it is optimal for the government to *not* differentiate the level of activity between aware and non-aware households.[42] The law of motion thus is unchanged compared to (1) except for the modified effective infection rate, $\tilde{\beta}$, and the same holds true for the government's program and first-order condition. Under the functional form Assumption 3, the optimal activity level satisfies condition (6), except that β is replaced by $\tilde{\beta}$.

Turning to laissez faire, aware households choose activity level $a^\star$ as they know from their test results that they no longer face infection risk. The average activity level in the economy thus is $\bar{a}(y) \equiv \sigma y a^\star + (1-\sigma y)a(y)$, and the law of motion changes to[43]

$$\frac{dy}{dt} = f(y, a(y)) \equiv g(\bar{a}(y))\beta\bar{y}y\left(1-\left(\frac{y}{\bar{y}}\right)^{\omega}\right). \tag{10}$$

This has two consequences for the program of unaware households, whose value function we denote by $U(y)$. First, these households face higher infection risk because $\bar{a}(y) > a(y)$ enters the law of motion (10). But second, they anticipate the possibility of a private capital gain, $U^\star - U(y)$, from being infected and receiving a positive test result with the consequence of being released in the aware pool.

Formally, under Assumptions 3 and 4, unaware households solve

$$\max_{a_i} u(a_i) + \frac{a_i}{a(y)} f(y, a(y))\left(-\zeta\psi + \sigma(U^\star - U(y))\right) \quad \text{s.t. (10)},$$

which yields the first-order condition

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

Imposing functional form Assumption 3, $a^\star = 1$, and the equilibrium requirement $a_i = a(y)$, we thus find

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

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

Note that the static and dynamic externalities encountered in the baseline model continue to be present. But information about infection status introduces a third effect, which runs counter to the dynamic externality: The capital gains that households experience when transitioning from the unaware to the aware pool parallels the social welfare gain due to increased immunity, $V'(y)$, in the

[37] In Garibaldi et al. (2024) static externalities disappear when the health consequences of contacts between infectious and susceptible households exhibit constant returns to scale.

[38] Garibaldi et al. (2024) interpret some of their results by comparing optimum and equilibrium activity levels in time space (see, e.g., the discussion of their figure 2 on p. 30). As mentioned earlier, such comparisons disguise differences in incentives that derive from the epidemiological state rather than time.

[39] Quarantine measures also help to allocate resources more efficiently by targeting interventions. Our model abstracts from targeting and assumes that test, trace, and quarantine policies have zero (or some exogenous) costs.

[40] More intense testing might require activity because individuals have to go to testing sites or purchase tests. In contrast, Alvarez et al. (2021) and Piguillem and Shi (2022) assume that testing is independent of activity.

[41] In Gonzalez-Eiras and Niepelt (2020b) and Gonzalez-Eiras and Niepelt (2023), we allow for permanent separation between the infectious group and the rest of the population. Fetzer and Graeber (2021) analyze the effectiveness of a UK test-and-trace regime, finding that the policy significantly lowered infection rates.

[42] Suppose the activity level A of aware households exceeds the level a of unaware households, such that $\bar{a} = (\sigma y A + (1-\sigma y)a) > a$. Relative to a common activity level a, this yields the gain $\sigma y(u(A) - u(a)) \leq \sigma y u'(a)(A-a)$ (due to concavity of u) and the cost $(\psi - V'(y))(g(\bar{a}) - g(a))\tilde{\beta}\bar{y}y\left(1-(y/\bar{y})^{\omega}\right) \geq (\psi - V'(y))g'(a)(A-a)\sigma y\tilde{\beta}\bar{y}y\left(1-(y/\bar{y})^{\omega}\right)$ (due to weak convexity of g). Since the first-order condition for a common activity level implies $u'(a) = (\psi - V'(y))g'(a)\tilde{\beta}\bar{y}y\left(1-(y/\bar{y})^{\omega}\right)$, the cost of a discriminatory policy weakly exceeds the gain. The reverse argument also rules out $A < a$ as an optimum.

[43] Unlike in the government program, the effective infection rate β is unaffected because aware households do not have incentives to quarantine.

government's optimality condition. This new effect as well as the changed effective infection rate in the government program imply modified expressions for the static and dynamic externalities, respectively:

$$\mathcal{S}(a,y) = \beta\bar{y}y\left(1-\left(\frac{y}{\bar{y}}\right)^{\omega}\right)\psi\left[\zeta\frac{g(\bar{a})}{a}-(1-\kappa\sigma)g'(a)\right],$$

$$\mathcal{D}(a,y) = \beta\bar{y}y\left(1-\left(\frac{y}{\bar{y}}\right)^{\omega}\right)\left[(1-\kappa\sigma)g'(a)V'(y)-\sigma\frac{U^{\star}-U(y)}{a}g(\bar{a})\right].$$

Compared to the baseline model, the static externality is less negative (and might be positive).[44] Likewise, the dynamic externality is less positive. In general, the effect of testing on the critical threshold y^c therefore is ambiguous.

The following proposition characterizes the effects of testing:

Proposition 3. *Under Assumptions 1 and 4 and if* $T=\infty$, $g(a)=a^n$, *and* $\zeta \le n$,

i. *The government's value function and optimal activity are increasing in testing effectiveness,* σ.
ii. *For* $y \downarrow 0$, *equilibrium activity is increasing in testing effectiveness.*

The result in part i. states that the effect of testing effectiveness is to increase welfare and, consistent with Proposition 1, optimal activity. Intuitively, a change of σ affects the government's first-order condition (5) only by lowering the effective infection rate $\tilde{\beta}$ (contained in f). More effective testing thus reduces the health benefit of depressing activity and therefore increases the optimal activity choice.

Part ii. states that at low infection levels, equilibrium activity also increases in σ. The probability of private capital gains due to infection-induced awareness provides strong incentives for unaware households to be more active when these capital gains are large (i.e., when $U(y)$ is low), which is the case relatively early in the epidemic.

Proposition 3 i. suggests that lockdowns and testing are substitutes when testing also reduces infections, for instance because it is connected with contact tracing and quarantining. In the quantitative analysis in section 6, we verify that this is indeed the case. This is consistent with numerical results in Piguillem and Shi (2022), according to which tests are a substitute for lockdowns. Similarly, Pollinger (2023) finds that testing allows to relax lockdowns. In that model, infections stop once the infectious pool falls below a critical mass.

6. Quantitative results

Exploiting the gains in computational efficiency afforded by our framework, we turn next to its quantitative implications. We allow for endogenous costs of infection, vary the activity-infection nexus and the activity smoothing motive, change key epidemiological parameters, introduce stochastic regime change and constraints on the government's use of policy instruments, and analyze the consequences of testing and heterogeneity.

6.1. 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. Since the comparative advantage of our framework lies with the infection dynamics in the early stage of an epidemic, we report the average activity level until y 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 $y(0)$, 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 $V(y(0)) \approx -127.1$. This welfare difference corresponds to a gain from optimal policy of around 2513 U.S. dollars per capita in present value.[45] The equivalent permanent reduction in consumption, either in equilibrium or under the optimal policy, ϕ^u and ϕ^v respectively, solves

$$\frac{1}{\rho}(\ln(a^{\star}(1-\phi^u))-a^{\star}+1) = -150.3,$$

$$\frac{1}{\rho}(\ln(a^{\star}(1-\phi^v))-a^{\star}+1) = -127.1.$$

This yields $\phi^u - \phi^v = 0.0032$, i.e., the lifetime consumption equivalent of optimal policy equals 0.32 percent. We summarize these findings as well as the key statistics of all other scenarios in Table 2.

[44] In the baseline model, the static externality is proportional to $\zeta g(a)/a - g'(a) < 0$, where the inequality follows from the convexity of g and $\zeta < 1$. Since $g(\bar{a}) > g(a)$ and $1-\kappa\sigma \in (0,1)$, the result follows.

[45] See C.2 for details on this calculation. Farboodi et al. (2021) report a gain of around 4565 dollars. Their model features herd immunity effects and a higher fatality rate, and it assumes that recovered households increase their activity to $a^{\star}$.

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.

Scenario	y^c	$\hat{a}$	$\phi^u - \phi^v$
Baseline model	0.1027	0.6238	0.0032
Endogenous costs: Congestion	0.1383	0.6228	0.0046
Endogenous costs: Learning	0.0787	0.6618	0.0025
Linear activity-infections nexus	0.0322	0.7340	0.0054
Intermediate activity-infection nexus	0.0601	0.6629	0.0036
Testing and heterogeneity	0.1356	0.6835	0.0054
Regime change: Testing	0.1532	0.5633	0.0046
Regime change: Recurrent waves	0.3275	0.5509	0.0086
Constraints on policy instruments	0.1188	0.6057	0.0021
Stronger curvature of u	0.0766	0.7853	0.0019
Higher fatality rate	0.1301	0.5310	0.0058
Lower arrival rate of a cure	0.0259	0.7995	0.0043

6.2. 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 are quadratic in the infection flow, capturing congestion effects for instance due to capacity constraints in the healthcare sector. 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 C.2; in section 6.8 we analyze how higher total costs affect optimal policy). In the second specification, we assume that unit infection costs decrease with the cumulative share of the infected population, 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. This is broadly consistent with findings in Garibaldi et al. (2024) who argue that medical congestion externalities induce the planner to tighten social distancing measures in the early phase of the epidemic. Our results are also consistent with Alvarez et al. (2021) who find that it is only because of congestion effects that their model predicts a lockdown under the baseline calibration. In contrast, we find learning effects to shorten the optimal lockdown duration and to reduce the welfare gains from policy intervention.

6.3. 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, $n = 1.5$.

For fixed $a \leq 1$, a lower curvature parameter $n > 1$ increases the marginal cost of activity perceived by households, which is proportional to $g(a)/a = a^{n-1}$. Accordingly, equilibrium in an economy with smaller n exhibits lower activity—households behave more cautiously than in the baseline scenario. For the government, which perceives marginal costs proportional to $g'(a) = na^{n-1}$, the comparative statics are less clear cut since $g'(a)$ may rise or fall with n, depending on the value of a. For $a \geq \exp(-0.5) \approx 0.61$, which constitutes the relevant range, $g'(a)$ is strictly increasing in n as long as $n \leq 2$; that is, unlike households, the government increases activity when n falls.[46]

Fig. 4 shows optimal activity for $n = 1, 1.5, 2$, where we use Proposition 1 to compute activity when $n = 1.5$. As anticipated in section 2.3, the optimal policy when $n = 1.5$ is closer to the optimal policy under $n = 2$ than to the optimal policy under $n = 1$. When $n = 1$, the optimal lockdown is substantially shorter and less strict 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.

To the best of our knowledge, these are new results in the literature, which has struggled to analyze the epidemiologically relevant case of intermediate returns to scale; see Hethcote (1989) and our discussion in section 3.1.

6.4. Testing and heterogeneity

We simulate outcomes under the testing policy analyzed in section 5 for different detection probabilities.[47] A low detection rate, $\sigma = 0.1$, is of limited consequence compared to the baseline specification. But a higher detection rate of $\sigma = 0.5$ amplifies the differences between laissez faire and optimal government intervention, as Fig. 5 illustrates. The figure displays optimal (solid) and

[46] For this argument, we abstract from second-order effects on the government's choice operating through induced changes in the value function, see equation (5).
[47] We set $\kappa = 1$.

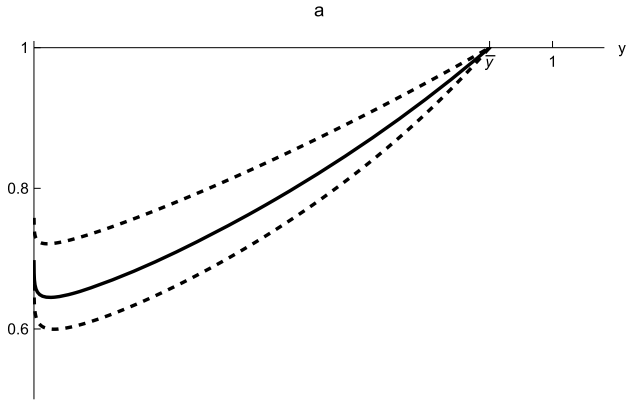

Fig. 4. Activity level in the government's program: $n = 1$ (dashed, top), $n = 1.5$ (solid), and $n = 2$ (dashed, bottom).

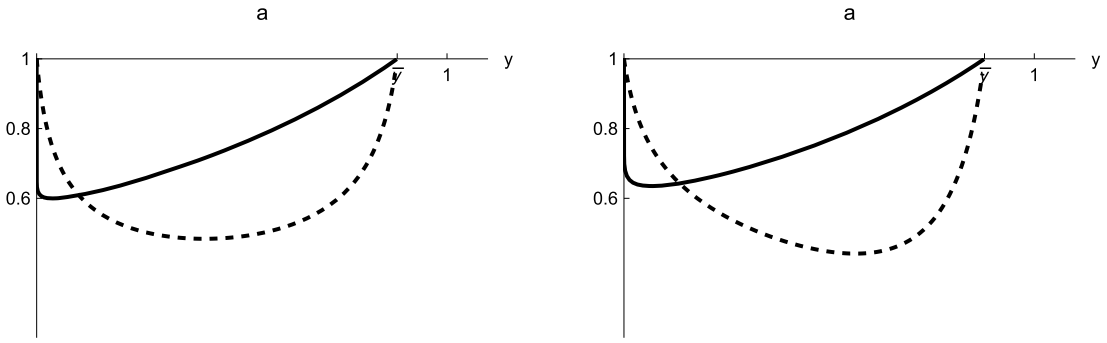

Fig. 5. Activity level in the government's program (solid) and in equilibrium (dashed): Baseline (left panel) and test, trace, and quarantine with $\sigma = 0.5$ (right panel).

equilibrium (dashed) activity in the baseline model (left panel) and with testing when $\sigma = 0.5$ (right panel). Modest effects on the government's preferred activity contrast with major changes in equilibrium activity.

Intuitively, unaware households behave less cautiously early in the epidemic because of the prospect of private capital gains upon becoming aware of their infection status, consistent with our third result in Proposition 3. In contrast, they behave less cautiously subsequently when the risk of meeting aware, infectious households has increased. As a consequence, lockdowns last longer than in the baseline, and the welfare gains from government intervention are larger. We emphasize that our results concern duration in state space, not time space; see the discussion in subsection 2.1.1.[48]

In line with our first result in Proposition 3, a marginal increase of σ evaluated at $\sigma = 0$ raises $V(y_0)$. Simulation results also show that the marginal welfare benefit of σ, $\partial V(y_0)/\partial\sigma$, is hump-shaped and approaches zero as σ increases and approaches one. This is consistent with the notion that convex costs of improving the efficacy of testing policies, as assumed in Alvarez et al. (2021), Piguillem and Shi (2022), or Pollinger (2023), would give rise to an interior optimum σ.

6.5. Stochastic regime changes

Next, we ask how stochastic change in the epidemiological environment affects optimal policy. We consider two scenarios. In the first, we allow for efficient testing to materialize stochastically, with arrival rate $\mu = 1/90$.[49] 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 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.[50]

The upside risk of effective testing in the first scenario increases the pre-regime-change values V and U, and it leads the government to wait longer before relaxing a harsher lockdown. This is consistent with results in Alvarez et al. (2021) according to which the anticipation of future testing results in an immediate lockdown, even when such a lockdown would not be imposed otherwise. It is

[48] The literature contains results according to which lockdowns last for fewer periods when they are accompanied by test, trace, and quarantine policies (Eichenbaum et al., 2021; Piguillem and Shi, 2022; Pollinger, 2023). These results need not contradict ours.

[49] This 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. Fetzer and Graeber (2021) show that this new regime lowered infection rates.

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

also consistent with Eichenbaum et al. (2022) who emphasize synergies between social distancing policies and testing/quarantining that arise because the former buy time for "testing/quarantining to come to the rescue" later (p. 2). The prospect of recurrent 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 relative to the baseline.

6.6. 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 stimulate it. 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.[51]

Ruling out inverse lockdown measures implies a longer and slightly stricter lockdown.[52] Although households act too cautiously during the later stage of the epidemic and the government cannot correct this, its own behavior during the early stage becomes more cautious as well: The government acts in this way to increase the likelihood that a cure will be found by the time it surrenders control over activity.

Importantly, the welfare gains of the optimal intervention are substantially reduced relative to the baseline, and this is even more the case when $n = 1$ (not reported), consistent with the discussion in section 6.3. 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 roughly 80 percent when $n = 1$. To the best of our knowledge, these are novel results in the literature.

6.7. Stronger curvature of u

Next, we consider a modified specification of preferences, letting net utility 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 of consumption equals one half rather than one. This reflects, for example, limited possibilities for households to self insure, rendering it costlier to lower activity. We stress that the literature avoids deviations from logarithmic or linear preferences, due to the resulting numerical complications. In contrast, our setup can easily deal with this generalization.

We find that, as a consequence of the lower elasticity of substitution, activity 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 optimal government intervention are smaller.

6.8. 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, and of a lower arrival rate of a cure that increases the expected duration to cure discovery from 1.5 to 5 years. In the first scenario, the government acts more cautiously while the opposite happens in the second. Intuitively, a smaller ν value limits the option to avoid future infections by delaying, and this reduces the incentive to sacrifice livelihoods in order to protect lives. 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 both scenarios.

6.9. Summary

We draw several conclusions from our quantitative analyses. First, the baseline scenario implies an optimal activity reduction during the early phase of the lockdown of roughly 38 percent.

Second, many of the qualitative comparative statics results are as expected: The optimal lockdown lengthens (in state space), and activity falls, when the fatality rate or the arrival rate of a cure increase, when congestion effects are stronger, or when learning effects are weaker. Testing policies that reduce the effective infection rate, or the prospect of such policies, increase the lockdown duration as well.

Third, other comparative statics results, including some novel ones, are more surprising, specifically in terms of size: The threat of recurrent infection waves triggers a large increase in the optimal lockdown duration and strictness, while a linear or near-linear activity-infection nexus or strong consumption smoothing needs have the opposite effect. The model implied welfare gains suggest that government intervention is particularly important when recurrent infection waves are a threat.

Finally, the most striking result concerns inverse lockdowns—stimulus measures that the government imposes once the economy starts to see the end of the tunnel: It is these measures that are responsible for a large, maybe even the dominant share of the welfare gains from optimal government intervention.

[51] At $\hat{y}$ say the government ends the lockdown (and would like to revert but cannot). We first solve for U and then find V and $\hat{y}$ from the conditions $V(\hat{y}) = U(\hat{y})$, $V'(\hat{y}) = U'(\hat{y})$.

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

Table 3
Temporary reductions in $g(a)$ in the SIR model such that z_∞ falls by one percent relative to the outcome with unit activity.

$g(a)$	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

7. Conclusion

We have developed a computationally efficient and flexible model of epidemic control and equilibrium dynamics. Households adjust 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 theoretical results whose intuitions extend to frameworks with larger state spaces. Moreover, we have exploited the gain in computational efficiency afforded by our model to numerically characterize equilibria and optimal government interventions in a large set of scenarios. Our findings may be summarized as follows:

First, a lockdown—government measures to curtail activity below its equilibrium level—is followed by the opposite, an inverse lockdown that stimulates activity. The reason for this is that, eventually, activity generates positive externalities because it drives the economy out of the crisis range with high infections. A substantial part of the welfare gains of optimal government intervention stems from the inverse lockdown. Second, testing policies may substitute for lockdowns but lockdowns are stricter when testing capabilities are expected to improve in the future.

Third, the policy implications of the baseline model are robust to many variations in parameter values and scenarios. Calibrated to match features of the COVID-19 pandemic in the U.S., the model's baseline specification suggests that activity should immediately have been reduced by 38 percent in mid March 2020 to yield welfare gains of 0.32 percent of lifetime consumption, equivalent to roughly 2500 U.S. dollars per capita in present value. Interestingly, a large part of these welfare gains arise from inverse lockdowns.

Fourth, there are three scenarios that call for major changes of policy relative to the baseline: When activity smoothing is essential because households cannot self insure their consumption, or when the activity-infection nexus is linear or near-linear, shorter and more lenient lockdowns are needed. In contrast, the threat of recurrent waves calls for longer and stricter lockdowns.

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.[53] The results on lockdowns and inverse lockdowns also call for a richer analysis of the fiscal consequences of optimal policy.

Scientists agree that humanity will be confronted with another pandemic in the not too far future. As in the spring of 2020, the key policy decisions to address the difficult tradeoffs between lives and livelihoods will have to be taken in the first few weeks of that pandemic. Our model captures these tradeoffs well, in particular if the likelihood of a cure is relatively high, and its flexibility and computational efficiency allows to identify robust policy responses.

Appendix A. Limitations of the one-state-variable approach

Since the long-run share of susceptible households is endogenous in the SIR model but exogenous in the logistic framework, some policy tradeoffs are present in the former but absent in the latter model. The question is how relevant these tradeoffs are during the early stage of an epidemic. To answer this question we pursue three approaches; all imply that the tradeoffs are not relevant during the early stage.

First, we find the duration D such that reducing activity from 1 to a at the beginning of the epidemic for the duration D, and reverting back to unity thereafter reduces z_∞, the long-run share of the recovered population in the SIR model, by one percent. Table 3 reports how D as well as the share of the population that undergoes infection while activity is reduced depends on the value of a (indexed by $g(a)$ for generality). For example, if activity is reduced by little, $g(a) = 0.9000$, then z_∞ falls by one percent when the activity reduction lasts for $D = 139$ days; the share of the population that undergoes infection during this period exceeds twenty percent. If $g(a) = 0.4250$, in contrast, then z_∞ falls by one percent when activity is reduced over a very long period, $D = 3358$ days, and the share of the population that undergoes infection in this case equals just 3.49 percent. (Even stronger activity reductions are ruled out.[54])

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—rendering herd immunity considerations unimportant. When we double the basic reproduction number then the critical y value is even higher, roughly 15 percent.

[53] 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.

[54] For $g(a) < 0.4168$ the infected pool shrinks during the intervention. Pollinger (2023) studies optimal suppression policies. For $g(a) > 0.9768$ the effect on z_∞ is smaller than one percent even if $D \to \infty$.

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.[55] 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.[56] Associating y in the logistic model with $\imath + z$ in the SIR simulation, we 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 Fig. 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.

Third, we feed the optimal policy during the first two years of the epidemic for another scenario, as reported by Farboodi et al. (2021) in their figure 4. We stipulate plausible continuation policies after the first two years that differ strongly, but we find that welfare under the different continuation policies is essentially unchanged (and differs by less than one percent from the reported number).[57]

The fact that our model abstracts from policy implications for the long-run share of susceptible households implies that it predicts optimal activity close to one when motives to delay infections are absent (i.e., when $\rho + \nu \to 0$). Even in this case, however, the predictions of our model and of SIR-based analyses remain very similar during the early phase of an epidemic because, as argued above, SIR-based analyses also abstract from these implications during the early phase.

Appendix B. Proofs

B.1. Proof of Lemma 1

Proof. When $T < \infty$, Assumptions 1 and 2 imply $V(y,0) = U^\star$ 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, $V(y,d) < U^\star$ for all $(y,d) \in (0,\bar{y}) \times (0,T]$ since an epidemic involves costs. When $T = \infty$, parallel logic implies $V(0) = V(\bar{y}) = U^\star$ and $V(y) < U^\star$ 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 $P^\star$ 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 $\omega, u'(\underline{a})$) 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 $V(0) = V(\bar{y}) = U^\star$, 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 $P^\star$ 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). □

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

$$(\rho + \nu)V(y,d) = u(a(y,d)) - u'(a(y,d))\frac{g(a(y,d))}{g'(a(y,d))} - V_d(y,d) + \nu U^\star \tag{11}$$

[55] Farboodi 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).

[56] To recover the numerical values we use the software WebPlotDigitizer available at https://apps.automeris.io/wpd/.

[57] One continuation policy features constant activity at the last reported value in figure 4; another policy adjusts activity to maintain the effective reproduction rate at the last reported value.

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 (11) 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 (11), 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 (11) 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 $[y_0, y_1), [y_1, y_2), \ldots, [y_{n-1}, y_n]$ where $y_0 = y_a$, $y_n = y_b$ and $\epsilon \equiv \sup_i \{y_i - y_{i-1}\}_{i=1,2,\ldots,n}$ is arbitrarily close to zero. Recall that a is non-monotone in each subinterval i.[58] Denote by $t(y_i)$ the time at which the economy reaches state y_i and let $\Delta t(\epsilon) = \sup_i \{t(y_i) - t(y_{i-1})\}_{i=1,\ldots,n}$. 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 $\{t(y_i)\}_{i=1,\ldots,n}$ as the original policy.[59] 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 $\{t(y_i)\}_{i=1,\ldots,n}$. 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.

Assume for notational simplicity that the horizon is infinite. Since the value function is differentiable and policy continuous; and since $U^\star = 0$ and $V(0) = 0$, equation (7) implies that $\lim_{y \downarrow 0} u(a(y)) - u'(a(y)) \frac{g(a(y))}{g'(a(y))} = 0$. Since for $a \in [0, a^\star]$ we have $u(a) \leq 0$ and $u'(a) \frac{g(a)}{g'(a)} \geq 0$, both with equality only if $a = a^\star$, it follows that $\lim_{y \downarrow 0} a(y) = a^\star$.

Finally, consider the alteration of a parameter other than ρ, ν, or T (which determines d). If this alteration changes $V(y, d)$ then, from relation (7), it must also change $\phi(a(y, d)) \equiv u(a(y, d)) - u'(a(y, d)) \frac{g(a(y,d))}{g'(a(y,d))}$, and both changes must have the same sign. Since u is increasing in $[\underline{a}, a^\star]$ and strictly concave, and g strictly increasing and weakly convex, we have

$$\begin{aligned} \frac{d\phi(a)}{da} &= u'(a) - u''(a)\frac{g(a)}{g'(a)} - u'(a) + u'(a)\frac{g(a)}{g'(a)^2} g''(a) \\ &= -u''(a)\frac{g(a)}{g'(a)} + u'(a)\frac{g(a)}{g'(a)^2} g''(a) > 0. \end{aligned}$$

For $\operatorname{sgn}(\Delta V(y, d))$ to equal $\operatorname{sgn}(\Delta \phi(a(y, d)))$, we thus must have $\operatorname{sgn}(\Delta V(y, d)) = \operatorname{sgn}(\Delta a(y, d))$. □

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 (in the online appendix), which states that V is strictly convex in $[y^{\min}, \bar{y}]$ and

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

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. □

B.4. Proof of Proposition 3

Proof. Let $\tilde{f}(y, a; \sigma)$ denote $f(y, a)$ in (1) with β replaced by $\tilde{\beta} = (1 - \kappa\sigma)\beta$. By optimality, the total derivative of V with respect to σ exceeds the direct effect of a change in σ keeping the activity path constant. From the definition of V in equation (4) with f replaced by $\tilde{f}$, and from $\frac{d\tilde{f}(y,a;\sigma)}{d\sigma} = -\kappa \tilde{f}(y, a; 0)$,

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

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

$$\frac{\partial V(y)}{\partial \sigma}|_{\mathbf{a}} = \kappa \int_0^\infty \psi\ \tilde{f}(y,a;0)\ \mathrm{e}^{-(\rho+\nu)t} dt \geq 0.$$

Proposition 1 then implies that optimal activity increases in σ. This completes the proof of part i.

Part ii. follows from totally differentiating the first-order condition (5) (with $f(y,a)$ replaced by $\tilde{f}(y,a;\sigma)$) with respect to σ and a, holding $V'(y)$ fixed.

For part ii., consider the optimality condition characterizing the equilibrium activity choice $a_i(y)$,

$$u'(a_i(y)) = \frac{\zeta\psi - \sigma(U^\star - U(y))}{a_i(y)} g(\bar{a}(y))h(y),$$

where $\bar{a}(y) = \sigma y a^\star + (1-\sigma y)a_i(y)$ and $h(y) \equiv \beta\bar{y}y\left(1-\left(\frac{y}{\bar{y}}\right)^\omega\right)$. Totally differentiating with respect to $a_i(y)$ and σ yields

$$\left(u''(a_i(y)) - \frac{\zeta\psi - \sigma(U^\star - U(y))}{a_i(y)} g(\bar{a}(y))h(y)\left[\frac{g'(\bar{a}(y))}{g(\bar{a}(y))}(1-\sigma y) - \frac{1}{a_i(y)}\right]\right) da_i(y)$$
$$= \left(\frac{\zeta\psi - \sigma(U^\star - U(y))}{a_i(y)} g'(\bar{a}(y))(a^\star - a_i(y))y - \frac{(U^\star - U(y))}{a_i(y)} g(\bar{a}(y))\right) h(y) d\sigma.$$

The term in parentheses on the left-hand side is negative for $y \downarrow 0$. This follows from $u''(a_i(y)) \leq 0$, $\zeta\psi - \sigma(U^\star - U(y)) \geq 0$ (implied by the first-order condition and $g(\bar{a}(y))h(y) \geq 0$), and $g'(a_i(y))a_i(y) \geq g(a_i(y))$ (implied by weak convexity of g).

Our goal is to establish a condition under which the term in parentheses on the right-hand side is negative for $y \downarrow 0$, such that $da_i(y)/d\sigma > 0$ for $y \downarrow 0$. Note first that

$$U(y) \leq \frac{u(a(y)) - u'(a(y))\frac{g(a(y))}{g'(a(y))}}{(\rho+\nu)},$$

where $a(y)$ denotes the optimal policy; this follows from $V(y) \geq U(y)$ and the representation of $V(y)$ given in Proposition 1. The term in parentheses on the right-hand side thus satisfies

$$\frac{\zeta\psi - \sigma(U^\star - U(y))}{a_i(y)} g'(\bar{a}(y))(a^\star - a_i(y))y - \frac{(U^\star - U(y))}{a_i(y)} g(\bar{a}(y))$$
$$= \frac{u'(a_i(y))}{g(\bar{a}(y))h(y)} g'(\bar{a}(y))(a^\star - a_i(y))y - \frac{(U^\star - U(y))}{a_i(y)} g(\bar{a}(y))$$
$$\leq \frac{u'(a_i(y))}{g(\bar{a}(y))h(y)} g'(\bar{a}(y))(a^\star - a_i(y))y - \frac{U^\star}{a_i(y)} g(\bar{a}(y)) + \frac{u(a(y)) - u'(a(y))\frac{g(a(y))}{g'(a(y))}}{(\rho+\nu)a_i(y)} g(\bar{a}(y)).$$

Since $U^\star = 0$ and $u(a(y)) \leq 0$, it suffices to show that

$$\frac{u'(a_i(y))}{g(\bar{a}(y))} g'(\bar{a}(y))(a^\star - a_i(y)) - u'(a(y))\frac{g(a(y))}{g'(a(y))}\frac{g(\bar{a}(y))}{(\rho+\nu)a_i(y)}\frac{h(y)}{y} \leq 0.$$

Since $a(y) \leq a_i(y)$ for $y \downarrow 0$ (from Proposition 2), we have $u'(a_i(y)) \leq u'(a(y))$ for $y \downarrow 0$, and it therefore suffices to establish that

$$\frac{g'(\bar{a}(y))}{g(\bar{a}(y))}(a^\star - a_i(y)) \leq \frac{g(a(y))}{g'(a(y))}\frac{g(\bar{a}(y))}{(\rho+\nu)a_i(y)}\frac{h(y)}{y}.$$

Since the functions $g, g', h/y$ are bounded over the relevant domain, and since $a_i(y) \uparrow a^\star$ for $y \downarrow 0$, this condition is indeed satisfied for $y \downarrow 0$. We conclude that $\frac{da_i(y)}{d\sigma} > 0$ for $y \downarrow 0$. □

Appendix. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jedc.2025.105145.

References

Abel, A.B., Panageas, S., 2020. Optimal management of a pandemic in the short run and the long run. Working Paper 27742. NBER, Cambridge, Massachusetts.

Acemoglu, D., 2009. Introduction to Modern Economic Growth. Princeton University Press, Princeton.

Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D., 2021. Optimal targeted lockdowns in a multigroup SIR model. Am. Econ. Rev. Insights 3 (4), 487–502.

Alvarez, F., Argente, D., Lippi, F., 2021. A simple planning problem for COVID-19 lockdown. Am. Econ. Rev. 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., Capuzzo-Dolcetta, I., 1997. Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations. Springer Science and Business, New York.

Bartsch, S.M., Ferguson, M.C., McKinnell, J.A., O'Shea, K.J., Wedlock, P.T., Siegmund, S.S., Lee, B.Y., 2020. The potential health care costs and resource use associated with COVID-19 in the United States. Health Aff. 39 (6), 927–935.

Bethune, Z.A., Korinek, A., 2020. Covid-19 infection externalities: Trading off lives vs. livelihoods. Working Paper 27009. NBER, Cambridge, Massachusetts.
Bisin, A., Moro, A., 2022. Learning epidemiology by doing: the empirical implications of a Spatial-SIR model with behavioral responses. J. Urban Econ. 127, 103368.
Calvia, A., Gozzi, F., Lippi, F., Zanco, G., 2023. A simple planning problem for COVID-19 lockdown: a dynamic programming approach. Econ. Theory. https://doi.org/10.1007/s00199-023-01493-1.
Crandall, M.G., Lions, P.-L., 1983. Viscosity solutions of Hamilton-Jacobi equations. Trans. Am. Math. Soc. 277 (1), 1–42.
Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2021. The macroeconomics of epidemics. Rev. Financ. Stud. 34 (11), 5149–5187.
Eichenbaum, M.S., Rebelo, S., Trabandt, M., 2022. The macroeconomics of testing and quarantining. J. Econ. Dyn. Control 138, 104337.
Ellison, G., 2024. Implications of heterogeneous SIR models for analyses of COVID-19. Rev. Econ. Des. https://doi.org/10.1007/s10058-024-00355-z.
Fajgelbaum, P.D., Khandelwal, A., Kim, W., Mantovani, C., Schaal, E., 2021. Optimal lockdown in a commuting network. Am. Econ. Rev. Insights 3 (4), 503–522.
Farboodi, M., Jarosch, G., Shimer, R., 2021. Internal and external effects of social distancing in a pandemic. J. Econ. Theory 196 (C), 105293.
Ferguson, N.M., Laydon, D., Nedjati-Gilani, G., collaborators, 2020. Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Report 9. Imperial College, London.
Fetzer, T., 2022. Subsidising the spread of COVID-19: evidence from the UK's eat-out-to-help-out scheme. Econ. J. 132 (643), 1200–1217.
Fetzer, T., Graeber, T., 2021. Measuring the scientific effectiveness of contact tracing: evidence from a natural experiment. Proc. Natl. Acad. Sci. USA 118 (33). https://doi.org/10.1073/pnas.2100814118.
Garibaldi, P., Moen, E.R., Pissarides, C.A., 2024. Static and dynamic inefficiencies in an optimizing model of epidemics. Econ. Theory 77 (1), 9–48.
Gersovitz, M., Hammer, J.S., 2004. The economical control of infectious diseases. Econ. J. 114, 1–27.
Giannitsarou, C., Kissler, S., Toxvaerd, F., 2021. Waning immunity and the second wave: some projections for SARS-CoV-2. Am. Econ. Rev. Insights 3 (3), 321–338.
Gonzalez-Eiras, M., Niepelt, D., 2020a. On the optimal 'lockdown' during an epidemic. Covid Econ. 7, 68–87.
Gonzalez-Eiras, M., Niepelt, D., 2020b. Optimally controlling an epidemic. Discussion Paper 15541. CEPR.
Gonzalez-Eiras, M., Niepelt, D., 2020c. Tractable epidemiological models for economic analysis. Discussion Paper 14791. CEPR.
Gonzalez-Eiras, M., Niepelt, D., 2022. The political economy of early COVID-19 interventions in US states. J. Econ. Dyn. Control 140, 104309.
Gonzalez-Eiras, M., Niepelt, D., 2023. Optimal epidemic control. Discussion Paper 23-11. University of Bern.
Hall, R.E., Jones, C.I., Klenow, P.J., 2020. Trading off consumption and COVID-19 deaths. Q. Rev. - Fed. Reserve Bank Minneap. 42 (1), 2–13.
Hethcote, H.W., 1989. Three basic epidemiological models. In: Gross, L., Hallam, T.G., Levin, S.A. (Eds.), Applied Mathematical Ecology. Springer, Berlin, pp. 119–144.
Hethcote, H.W., 2000. The mathematics of infectious diseases. SIAM Rev. 42 (4), 599–653.
Jones, C.J., Philippon, T., Venkateswaran, V., 2021. Optimal mitigation policies in a pandemic: social distancing and working from home. Rev. Financ. Stud. 34 (11), 5188–5223.
Kaplan, G., Moll, B., 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., McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. Proc. R. Soc., Ser. A 115 (772), 700–721.
Kolmogorov, A.N., Fomin, S.V., 1970. Introductory Real Analysis. Prentice Hall, New Jersey.
May, R.M., 1974. Biological populations with nonoverlapping generations: stable points, stable cycles, and chaos. Science 186 (4164), 645–647.
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., Halverson, P.K., 2020. Population point prevalence of SARS-CoV-2 infection based on a statewide random sample—Indiana, April 25–29. Morb. Mort. Wkly. Rep. 69, 960–964.
Miclo, L., Spiro, D., Weibull, J.W., 2022. Optimal epidemic suppression under an ICU constraint: an analytical solution. J. Math. Econ. 101, 102669.
Piguillem, F., Shi, L., 2022. Optimal Covid-19 quarantine and testing policies. Econ. J. 132 (647), 2534–2562.
Pollinger, S., 2023. Optimal contact tracing and social distancing policies to suppress a new infectious disease. Econ. J. 133 (654), 2483–2503.
Rachel, L., 2022. An Analytical Model of Behavior and Policy in an Epidemic. Unpublished. London School of Economics.
Richards, F.J., 1959. A flexible growth function for empirical use. J. Exp. Bot. 10 (2), 290–301.
Toxvaerd, F., 2020. Equilibrium social distancing. Working Paper 2020/08. Cambridge Institute for New Economic Thinking, Cambridge.
Wu, K., Darcet, D., Wang, Q., Sornette, D., 2020. Generalized logistic growth modeling of the COVID-19 outbreak: comparing the dynamics in the 29 provinces in China and in the rest of the world. Nonlinear Dyn. 101 (3), 1561–1581.