Optimal COVID-19 Quarantine and Testing Policies*

Facundo Piguillem† Liyan Shi‡

December 2021

Abstract

We study quantitatively the optimality of quarantine and testing policies; and whether they are complements or substitutes. We extend the epidemiological SEIR model incorporating an information friction. Our main finding is that testing is a cost-efficient substitute for lockdowns, rendering them almost unnecessary. By identifying carriers, testing contains the spread of the virus without reducing output. Although the implementation requires widespread massive testing. As a byproduct, we show that two distinct optimal lockdown policy types arise: suppression, intended to eliminate the virus, and mitigation, concerned about flattening the curve. The choice between them is determined by a "hope for the cure" effect, arising due to either an expected vaccine or the belief that the virus can be eliminated. Conditional on the policy type, the intensity and duration are invariant to the welfare function's shape: they depend mostly on the virus dynamics.

JEL classification: E1, E65, H12, I1

Keywords: COVID-19. Optimal quarantine. Optimal testing. Welfare cost of quarantines.

*We thank Luigi Guiso, Claudio Michelacci, Daniele Terlizzese, and seminar participants at PSE, Bocconi, EUI, and the Graz Schumpeter Center Covid-19 Workshop for providing us with thoughtful suggestions. Ruben Piazzesi provided superb research assistance. All remaining errors are our own.

†EIEF and CEPR (e-mail: facundo.piguillem@gmail.com).

‡EIEF and CEPR (e-mail: liyan.shi@eief.it).

1 Introduction

The COVID-19 pandemic took the world by surprise, driving many governments to take preventive measures. As a result, they imposed "lockdown" policies with heterogeneous degrees of intensity. These policies reduce or cease certain activities, with a high cost in terms of foregone production. The key element giving rise to lockdowns is the governments' inability to identify the virus carriers. If the problem is the lack of information, wouldn't it be more efficient to spend resources on gathering information and avoid lockdowns? This paper studies how testing can substitute for or complement lockdowns. We analyze the optimal joint time paths of quarantine and testing policies and argue that, although imprecise, widespread random testing policies can eliminate the need for lockdowns.

To analyze this problem, we build on the SEIR framework by Atkeson (2020).[1] In this environment, there is an outbreak of an infectious disease, which spreads through interactions between virus carriers and susceptible subjects. We extend the setting in several aspects. First, to model the *information friction*, we assume that exposed individuals are initially asymptomatic carriers who can also transmit it. Second, we introduce the possibility of a *critical mass* such that, if the number of carriers is above it, the virus reproduces; otherwise, it vanishes. To reflect that hospitals can be overwhelmed, we impose a *healthcare capacity constraint* limiting the number of patients that can be treated at a given time. An individual requiring medical care dies with a higher probability without proper care. Moreover, following Eichenbaum et al. (2020a), we incorporate *endogenous social distancing* measures undertaken by each individual besides the government policies. As the population learns about the virus spread, they restrain from engaging in social and economic activities.

The information friction plays a fundamental role in shaping containment measures. If the infectious state were fully observable, containing the virus would not be a problem. A social planner would simply isolate (quarantine) all infected individuals until the illness recedes while letting the unaffected population unrestricted. Without information, the government can only force "indiscriminate" quarantines: both the infected and the healthy must be isolated. Thus, information has value and can be obtained by testing. However, the testing technology is imperfect, flagging as infected some people who are not (false positives) and flagging as uninfected some who carry the virus (false negatives). Since testing is costly and provides imprecise information, its usefulness is not guaranteed.

The critical mass is specific to our setting and may initially appear arbitrary, so further explanations are in order. It is a sensible assumption that rather than being a driver of

[1] In two contemporaneous papers, Eichenbaum et al. (2020a) and Alvarez et al. (2021) rely on the classical SIR model, which does not distinguish between symptomatic and asymptomatic individuals.

the results, it makes them stronger: with or without the critical mass, widespread random testing dominates lockdowns. The SIR paradigm assumes divisible (continuous) humans. In general this assumption is innocuous and renders tractable a complex problem. However, the subjects' indivisibility, together with proportional reproduction rates, imply that the virus cannot be eliminated. No matter how small the "infected mass" is, it is never zero, and it eventually grows back. Hence, suppression strategies directed to eliminate the virus irrevocably are (implicitly) assumed futile. Nevertheless, many countries, perhaps based on the success story after the 2002 SARS-CoV-1 crisis, tried suppression.[2] Given this previous success and that suppression may render testing obsolete, it is nothing but fair to give it a chance. Last but not least, we also show that the expectation of the potential arrival of a medical innovation generates analogous effects to the critical mass. We term this effect the "hope-for-the-cure" induced suppression.

We analyze the problem quantitatively and in sequential order. We first study the optimal intervention without testing and then the optimal joint paths of quarantines and testing. We start characterizing how much and for how long activities should be restricted. Using this as a benchmark, we then analyze how testing changes the lockdown patterns and how it is used, arguing that testing can substitute and is more efficient than lockdowns. Finally, we incorporate the expectation that a medical innovation could arrive. We show that for many relevant scenarios, this expectation leaves the results unaffected, although it can generate sudden changes in the types of lockdowns.

To quantitatively assess the policies, we calibrate the model to the Italian COVID-19 outbreak. To deal with widespread underreporting of cases and fatalities, we target the fatalities path using the total excess deaths relative to previous years. To discipline the population's endogenous reaction, we use the cellphone movement index constructed by Durante et al. (2021). Finally, for the scenarios with a critical mass, we calibrate it to be one individual. This is consistent with our interpretation that the critical mass is a way to undo the human divisibility assumption. Nevertheless, we present detailed results varying every main component determining the optimal intervention.

First, suppose that testing is either not possible or not used, then three types of optimal policies arise: **suppression**, **mitigation**, and **no intervention**. The preferred policy type depends on the aversion to output variation, the statistical value of life, the critical mass size, and the vaccine's expected arrival. However, *conditional on following one of these strategies, the intensity and the duration of lockdown are barely affected by the properties of the welfare*

[2] The draconian measures imposed in Wuhan resembled those after the reaction to SARS-CoV-1. Also, during the initial outbreak, the government isolated 11 municipalities in Italy, even preventing the citizens from leaving their homes. Many researchers were wondering whether an approach as the one used with SARS-CoV-1 would work with COVID-19; see, for instance, Wilder-Smith et al. (2020).

function or the value of life.

In broad terms, when the value of life is below a threshold no intervention is optimal. Only individuals' self-imposed precautions control the spread. For higher life values intervening becomes optimal, but the chosen strategy depends on whether there is critical mass and the vaccine's expected arrival time, i.e., there is hope for a cure. When there is no "hope," the planner always appeals to mitigation strategies: it "flattens the curve" to avoid overwhelming the hospitals until herd immunity arrives. Instead, if there is hope for a cure, the planner may choose a suppression strategy consisting in a decisive intervention, shutting down a large share of activities until the virus is eliminated, either because it falls below the critical mass or a medical innovation arrives.[3]

It is worth mentioning that the expectation of a vaccine not always generates strategy switches; it does so only when its option value is large enough. But, the option value is not monotone. On one extreme, if the vaccine (or cure) is expected to arrive far enough in the future, it has minor consequences for current decisions with almost no effect on optimal policies. On the other extreme, if the vaccine is expected to arrive immediately, then the virus does not pose a threat, so its option value is also zero; thus, the planner never intervenes. For intermediate values of the arrival rate (in our quantitative results between 1 and 2 years), the option value is maximized and can generate relevant strategy jumps.

This finding rationalizes the observed heterogeneity in responses among countries. One could think that most of this heterogeneity is driven by the trade-off between lives and output cost. However, without the hope for the cure, only mitigation policies are optimal, which depends only on the virus dynamic. The different beliefs about the possibility of a cure drive the bulk of disparate responses, by generating jumps to suppression strategies. For instance, some countries that are smaller, or more isolated, or have better control over their borders may intent to suppress, hoping to hit the critical mass. In other countries, instead, even if the authorities believe that the critical mass is not reachable, they may be optimistic about the arrival of a vaccine. This optimism is not only linked to the medical discovery alone but also to early access and implementation. In any case, this heterogeneity is likely to be observed only on the outbreak's initial stages. As the information settle, beliefs should converge and so should the implemented policies.

These policies have significant output costs, which can fall by more than 60% at the intervention's peak. This brings about the possibility of complementing the quarantine with testing to simultaneously slow down the reproduction and avoid the output cost. By testing,

[3] Technically, when there is the possibility of ending the disease, the welfare function is convex. In our case, there are two local maxima, one corresponding to the optimal suppression and the other to optimal mitigation. Their relative level depends on the value of life, discounting, aversion to variation, etc. This also warns those using a first-order approach, which we do not.

the government can identify exposed individuals. Once identified as positive, a subject must endure a (personal) quarantine. Testing is done by *randomly selecting* individuals for whom there is no information yet: those who have never tested positive before. Identifying a positive case has two beneficial effects. First, it becomes possible to quarantine him/her, even in a stricter way than the rest of the population. Second, it identifies immunity: once recovered, the individual can return to work without restrictions. Given the different types of intervention, it is unclear under which strategy testing could be helpful, if any. For instance, if the government is seeking suppression, testing could be useless. Moreover, we consider only random testing, which is simple and easy to implement, but perhaps it is the least efficient. Thus, it could be discarded just on those grounds.

To evaluate the contribution of testing, we follow Atkeson et al. (2020) and calibrate its marginal cost to around $55 for relatively small testing intensities but increasing in the number of tests. The speed at which the marginal cost grows is chosen in such a way that it would be economically infeasible to test the entire population at once.[4]

We find that **testing is intensively used as a substitute for indiscriminate quarantines** and **generates substantial welfare gains**. The output gains are so large that lockdowns could be completely avoided, even though they are still used moderately. In our favorite mitigation scenario, testing is used intensively. On average, 20% of the unidentified population is tested every day the first month and then continues with lower intensity for about a year, averaging 8% of the population per day.[5] This policy is very costly, amounting to 1.3% of annual GDP. But this cost is easily compensated for by eliminating the lockdown. This strong substitutability is also present when the government follows a suppression strategy. With testing, instead of shutting down 60% of the economic activities, it shuts down only 17%, and instead of doing it for 80 days, the shutdown lasts for only three weeks. Here, the total cost is just 1.1% of annual GDP.

In short, whether the government chooses mitigation or suppression depends on a combination of factors, but either way, testing is used extensively, well beyond the magnitudes observed in most countries, and the lockdowns are either reduced or eliminated.[6]

[4] We do not know how much and how fast the testing capacity can be expanded, especially at the initial moment of an outbreak. We conjecture that if the population is sufficiently big, expanding the capacity on the spot to test the whole population would be infeasible.

[5] In Italy, the testing intensity barely reached 1 million people a day, while the optimal policy requires 12 million tests per day during the first month and 4.8 million after that.

[6] China is an exception in this regard. After experimenting with lockdowns as in Wuhan, when the city of Qingdao appeared to be affected in October 2020, the government decided to test the entire city.

1.1 Literature review

The literature on epidemiology control dates back to the framework proposed by Kermack and McKendrick (1927), also known as the SIR epidemiology model. Until COVID-19, the economic literature on epidemics was scarce, with some exceptions mostly related to HIV, as Greenwood et al. (2019).

The large wave of papers on COVID-19 was triggered by Atkeson (2020) and Eichenbaum et al. (2020a). Atkeson (2020) is a simple exercise to compute the disease's projected paths and evaluate its expected economic impact. We build on this work by deepening the information structure. The author assumes that only symptomatically infected subjects are contagious. We instead assume that the symptomatically infected are isolated and, thus, do not infect others. The asymptomatic individuals are the ones fueling the spread of the disease. This extension allows us (1) to better fit the dynamics of the disease and (2) to incorporate a well-defined information friction generating the need for testing.

The closest paper to ours is Alvarez et al. (2021), who study optimal interventions in a SIR framework. They assume a linear welfare function, weighting output and the statistical value of life, and solve an optimal control problem on the unrestricted quarantine policy space while constraining the "test-tracing."[7] We differ in several dimensions. We use a different meeting technology (theirs exhibits increasing returns, while ours is scale independent), different information structure (SEIR vs. SIR), alternative welfare functions (non-linear) and we consider the possibility of additional population's responses. These differences imply contrasting messages regarding testing. They find, in a quantitative example, that random testing would never be optimal, while we argue that even random testing is a better technology than indiscriminate quarantines. Their simplifying assumption allows them, though, to study the unrestricted space of policies while we restrict the space of policies and consider a simple form of endogenous population reaction.

The economic literature contributed to the epidemiological one by introducing non-mechanic responsive individuals. Eichenbaum et al. (2020a) were the first to point the relevance of the endogenous agents' reactions. They build a SIR model adding standard macroeconomic features. Even though agents take preventive measures, the competitive equilibrium is suboptimal because agents do not internalize the effect of their actions on others. They then analyze the optimal Pigouvian tax correcting the externality. Focusing on this externality, Farboodi et al. (2021) retain the SIR framework and put additional structure on the way in which interactions (matches) take place. Using location data, they are able to quantify the individual responses to the COVID-19 arrival and the optimal interven-

[7] They also explicitly consider the possibility that in some countries the lockdown could be less effective or harder to implement. For implications about lack of government commitment, see Moser and Yared (2020).

tion's value. Assenza et al. (2020) go deeper into characterizing social interactions as the outcome of a game, which provides a deep interpretation of the nature of the inefficiencies. None of these papers analyzes testing policies, which is our main focus. As Atkeson (2020), we study a SEIR rather than a SIR model.[8] As Eichenbaum et al. (2020a), we incorporate endogenous population's reactions, although without all their richness. As Farboodi et al. (2021), we use the estimations of Durante et al. (2021) to quantify the population's response.

While lockdowns attracted much attention from researchers, testing has been less studied. To the best of our knowledge, the first pointing out the substitutability between quarantines and testing was Dewatripont et al. (2020), who argue that testing was essential to "restart the economy." Later, with different degrees of complexity and focus, several authors develop this insight in some dimension or another. For instance, Berger et al. (2020) fixes the quarantine and testing policies, estimate a SEIR model, and show that testing is instrumental in softening the economic effects of the quarantine. Chari et al. (2021) introduce a noisy signal about the agent's infectious status and study targeted testing. They conclude that targeted testing is more cost-effective than mere isolation. Atkeson et al. (2020) emphasize the cost-effectiveness by comparing alternative scenarios using the, by then known, testing marginal costs. Eichenbaum et al. (2020b) incorporate testing into a version of Eichenbaum et al. (2020a) and study how alternative testing strategies affect individual behavior by changing their information sets.[9] Brotherhood et al. (2020) consider a richer age structure and endogenous meeting decisions and argue that testing dominates stay-at-home policies.[10] The main difference between our approach and the previous papers is that rather than comparing a few, even though relevant, scenarios, we do not take the policies as given, but we study the optimal joint paths of quarantine and testing. In this regard, through additional simplifying assumptions, Pollinger (2020) can provide some analytical characterizations.

We abstract from heterogeneity. Since our main findings are related to mass testing as substitutes for quarantines, including additional dimensions would not change the results. In any case, there is no doubt that heterogeneity is relevant for many questions. Among the many papers incorporating it and studying the distributional effects of policies are Brotherhood et al. (2020), Favero et al. (2020), Glover et al. (2020), and Kaplan et al. (2020).

[8]Bar-On et al. (2021) study the implications of using SEIR vs. SIR models. They show that the latent variable E is key to match the virus dynamics, with important implications for optimal policy.

[9]In their SIR model, agents do not know their health status. Hence, testing strategies change their information sets and therefore their behavior. In our paper, the government needs the information to design a fine-tuned quarantine policy. When an asymptomatic agent learns that she is infected, her response is restricted to zero interactions, while if symptomatic, both the agent and the government know it and act accordingly. In this sense, testing allows for "targeted lockdowns," as in Acemoglu et al. (2021).

[10]Their argument depends on the accuracy of the testing technology. We consider this fact in our calibration. For the relevance of antibody testing, see Guimarães (2021).

Figure 1: SEIR Model Transition

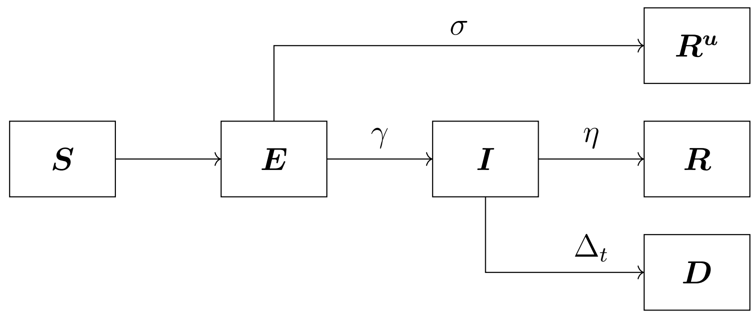

2 A SEIR model of disease contagion

Time is continuous and runs indefinitely, $t \in [0, \infty)$. At time 0, the economy is hit by a disease caused by a deadly virus that starts to spread. To capture the fast-moving dynamics, we use the convention that one unit of time is one day. The agents discount the future at rate $\rho > 0$. At each period t, there is population N_t, with an initial mass: $N_0 = 1$.

Each individual can be one of 5 types: susceptible, exposed, infected, identified recovered, or unidentified recovered. We denote by S the number of individuals still unaffected but susceptible to the virus. There are two types of contagious carriers: exposed asymptomatic E and infected symptomatic I. When an individual first contracts the virus, it always starts in group E. The asymptomatic E could become symptomatic I after an incubation period or may never show symptoms, in which case they remain in E until recovery.

When subjects recover, they become immune permanently.[11] Depending on the symptomatic history, an immune agent may not be identified. Let R denote the number of immune recovered agents who were previously symptomatically infected, and R^u the immune recovered subjects who never displayed symptoms. The former group has observable signals that make them identifiable, while the latter could remain unidentified without additional information. For this reason, the distinction between R and R^u is important. Figure 1 summarizes the transitions across different groups. It must be that:

$$N_t = S_t + E_t + I_t + R_t + R_t^u.$$

If the exposed population is above a **critical mass**, i.e., $E_t > \underline{E} \geq 0$, the virus spreads through meetings between exposed and susceptible individuals. Otherwise, it is

[11] There is consensus that most individuals who recover from COVID-19 would retain immunity for at least some time. Hansen et al. (2021) find that immunity would be around 80% after six months.

self-contained, and all patients gradually recover. This mass is specific to our setting, so some explanations are in order. It serves two purposes. 1) Without it, interventions aiming to eliminate the virus are pointless. Thus, successful strategies, as those used by China against SARS CoV-1, would be outright discarded. 2) We show in Section 5 that it has a "hope for the cure" effect akin to an early arriving vaccine.

Our SEIR framework assumes divisible (continuous) humans. In general, this assumption is innocuous and renders tractable a complex problem. However, subjects' indivisibility and proportional reproduction rates imply that the virus cannot be eliminated. No matter how small the "infected mass" is, it is never zero, and it eventually grows back. Hence, trying to eliminate the virus irrevocably would be impossible. Following this interpretation, we later calibrate $\underline{E}$ to be equivalent to one person. This makes sure that whenever less than an (entire) person is infected, the disease disappears.

To avoid the spread, the government can impose quarantines, either indiscriminate (lockdowns) or targeted. In a lockdown, the government simply shuts down a proportion q_t of *all* economic and social activities. Instead, with directed quarantines, the government singles out targeted groups and forces some or all of them to isolation. Targeted quarantines require information to identify the targeted groups. Since the symptomatic individuals I_t can be identified by their symptoms, it is optimal to force them into a full quarantine.[12] Hence, the virus spreads **only** through meetings between the exposed asymptomatic and unaffected individuals. Similarly, the recovered individuals R are known to be immune based on their symptomatic history and thus are not subject to restrictions.[13] In this section, we maintain the assumption that the government has no information about the identity of S_t, E_t or R_t^u individuals; hence, they are all subject to the lockdown q_t. In what follows we refer to interactions as economic activities, but the reader should bear in mind that our interpretation is also inclusive of social interactions that are not necessarily measured in the GDP.

Meetings are reduced due to the population's social distancing and government measures limiting human activities. The number of meetings could depend on how individuals respond to the risk implied by the virus and the restrictions imposed by the authorities. In addition, we assume that agents do not directly observe the death rate, so they must infer the death

[12]This assumption is in line with the preventive measures taken by all the governments as soon as they detect an positive case. However, there could be compliance issues. The infected could disregard government directives and still engage in social interactions, which seems common in the U.S. and U.K., see Atkeson et al. (2020) and Smith et al. (2021). Finally, the state of the spread may affect the optimality of quarantines for known infected cases. If, for instance, the mass of infected individuals is sufficiently high, or they are concentrated in strategic tasks, quarantining them could have significant adverse effects in production, with worse consequences than the virus.

[13]The introduction of a "green pass" in Europe has this feature. Vaccinated individuals are subject to few or no restrictions. Proof of recovery is considered equivalent to a complete vaccination cycle.

risk from the observed fatalities. To implement these ideas, we rely on a statistical model determining the fraction of meetings that each individual avoids after learning the number of fatalities and the restrictions imposed by the authorities. We denote by $b_t(q_t, D_t) \in [0,1]$ such fraction, representing the intensity of interactions. If $b_t = 1$, individuals avoid every type of social interaction, and thus, there are no meetings; if $b_t = 0$, individuals do not constraint themselves. For simplicity, in what follows, we skip the explicit dependency of b_t on its arguments whenever no confusion arises, but the reader should bear it in mind.

The number of infections depends on the number of individuals interacting adjusted by the intensity of the interaction: ι_t. Since identified recovered agents interact freely and the infected individuals are forced to isolate, the total measure of individuals interacting is:

$$\iota_t = (1 - b_t)(S_t + E_t + R_t^u) + R_t. \tag{1}$$

The voluntary reaction of the population is a key element to account for when designing the containment policy. There is ample evidence that in many countries, including Italy, the population reacted by taking precautionary measures when it was evident that COVID-19 was present and dangerous. This reduces the need for government intervention, and it could imply that no intervention is optimal.

Let $\lambda m(\iota_t, E_t(1 - b_t))$ be the meeting function between virus carriers and the rest of the population. Because the symptomatic are fully quarantined, the total number of virus spreaders is just E_t, hence the second term in the meeting function. Not all meetings generate an infection. Since only $\frac{S_t}{\iota_t}$ of workers are susceptible, and individuals are socially distanced, there are only $\lambda \frac{S_t(1-b_t)}{\iota_t} m(\iota_t, E_t(1 - b_t))$ meetings generating newly affected (exposed) individuals.[14]

Once exposed, an individual becomes symptomatically infected at rate γ and can recover at rate σ without ever being symptomatic. Thus, the exposed population follows the law of motion:

$$dE_t = \begin{cases} \left[\lambda \frac{S_t(1-b_t)}{\iota_t} m(\iota_t, E_t(1 - b_t)) - (\sigma + \gamma)E_t\right] dt, & \text{if } E_t \geq \underline{E} \\ -(\sigma + \gamma)E_t dt, & \text{if } E_t < \underline{E}. \end{cases} \tag{2}$$

The symptom's appearance and the asymptomatic's recovery rates, γ and σ, are independent of the state of the economy. They just reflect the individual's strength to fight the virus inside their biological system. The same is true for the contagion intensity rate λ, which is

[14] In Atkeson (2020), only the infected individuals can transmit the virus. Thus the infectious meetings are $m(\iota_t, I_t)$. The author does not distinguish between symptomatic and asymptomatic carriers, however. We borrow his notation and give it a different interpretation. In our setting, this distinction clarifies production effects and is instrumental when analyzing the information friction. Otherwise, we could merge them into the standard SIR model as a single type.

a scale parameter capturing the level of interactions among agents in their daily activities.

The speed at which the illness spreads is state-dependent, increasing in the number of exposed E_t and susceptible. The function $m(\iota_t, E_t(1-b_t))$ could incorporate potential "congestion" effects. For instance, one may think that, when most of the population is already affected, most meetings would be between immune and infected individuals and thus would generate fewer infections.

We emphasize the relevance of the critical mass $\underline{E}$ for policy interventions. If $\underline{E} > 0$, it is possible to take drastic measures to force the affected population below the critical mass so that the virus disappears and the infection is definitively defeated. In contrast, if $\underline{E} = 0$, the virus never dies out. Even when the exposed population is reduced to a negligible measure, the virus can always resurface and spread again.

As a result, the government could intervene with two alternative approaches, or types of strategies, that we term: **mitigation** and **suppression**. *Mitigation* is a strategy directed to slow down the spread until herd immunity. The policymaker only regulates the speed at which the number of exposed and infected subjects arrive; it "flattens the curve" until herd immunity puts an end to the problem. For this strategy, the spare hospital capacity is a first-order concern. In contrast, *suppression* is a strategy directed to stop the virus rather than slowing down the spread. It can be characterized by decisive interventions preventing a large proportion of interactions. The government aims to solve the problem without resorting to herd immunity; maybe eliminating the virus before is too widespread or due to an exogenous medical innovation.

The trade-off between these strategies depends on the effects of the virus on the population and the hospital capacity. We assume that the body's ability to dispose of the virus is independent of the health system. This reflects in a state independent recovery rate η. In contrast, the death rate could depend on the *capacity of the health system* to treat patients. Patients with severe cases could require medical assistance and potential hospitalization. However, in any given period t hospitals can only treat H_t patients. Once that capacity is exceeded, i.e., $I_t > H_t$, the treatment received by each patient is diluted. Those who are properly treated die at rate θ, while those untreated due to lack of capacity die at rate $\delta > \theta$. As a result, the average (state-dependent) daily death rate Δ_t for the infected agents satisfies:

$$\Delta_t = \theta \underbrace{\min\left\{1, \frac{H_t}{I_t}\right\}}_{\text{fraction treated}} + \delta \underbrace{\max\left\{1 - \frac{H_t}{I_t}, 0\right\}}_{\text{fraction untreated}}, \tag{3}$$

and the law of motion of infected individuals is:

$$dI_t = [\gamma E_t - (\eta + \Delta_t) I_t] dt. \tag{4}$$

The previous relations imply that the laws of motion for the recovered and the total population satisfy:

$$dR^u_t = \sigma E_t dt, \tag{5}$$

$$dR_t = \eta I_t dt, \tag{6}$$

$$dN_t = -\Delta_t I_t dt. \tag{7}$$

Indiscriminate quarantines q_t directly impact economic activity while social distancing is less damaging, reducing production in a proportion $o(b_t) \leq 1-q_t$. It is unclear how significant the effect of social distancing is on production. One can think that social distancing has no output cost. For instance, individuals still shop and buy the same value in goods, but they do it less often and with less direct physical contact, which would imply $o(b_t) = 0$. However, other economic activities could be affected. For example, agents could stop dining in restaurants, even when they remain open.

For simplicity, we assume a linear production technology in the active labor force: $Y_t = L_t$. Infected individuals cannot produce because they are sick and fully quarantined. Only the susceptible, the undetected exposed, and the fully recovered can produce. Indiscriminate quarantines prevent a fraction $q_t \in [0, 1]$ of the individuals with an unknown status from engaging in any economic activity. Then, the total production is:[15]

$$Y_t = (1 - q_t)(S_t + E_t + R^u_t) + R_t.$$

As a benchmark, the undistorted output would be $Y_t = S_t + E_t + R^u_t + R_t$. The only good produced is nonstorable, and there is no possibility of borrowing or saving. Hence, consumption must equal production in every period: $C_t = Y_t$.

2.1 Welfare evaluation

We now turn to the decisions of the policymakers. As a relevant benchmark, we start characterizing the optimal policy without the possibility of testing. Then, in Section 4, we show how testing alters and improves the results.

When intervening, the government considers the impact of its policies on the virus dynamics and the possibility of a future medical innovation. It could be a vaccine or an effective treatment. We follow Alvarez et al. (2021)'s and Shimer and Wu (2021)'s approaches. It consists in assuming that with Poisson arrival rate $\phi \geq 0$, a medical innovation arrives. Shimer and Wu (2021) assume that a perfect vaccine arrives, preventing any future spread but not curing the already infected. Alternatively, Alvarez et al. (2021) assume that a per-

[15] We allow the recovered whose status is observable to return to work, which is optimal.

fect cure arrives. As a result, the economy enters a new state without the disease upon the innovation's arrival. We denote this value with $V(I_t, E_t, N_t)$.

Formally, the policymaker chooses a path $\{q_t : t \geq 0\}$ to maximize social welfare,

$$\max_{\{q_t\}} \int_0^\infty e^{-(\rho+\phi)t} \left[u\left((1-q_t)(S_t + E_t + R_t^u) + R_t\right) - v\left(\Delta_t I_t\right) + \phi V(I_t, E_t, N_t)\right] dt, \tag{P}$$

subject to equations (1) to (6).

The function $u(\cdot)$ captures society's preferences for consumption. It must not be interpreted as just the private individuals' utility but as a composite encompassing potential externalities. Such possibility has been discussed in the literature, with a wide range of effects. Some authors, such as Assenza et al. (2020), Eichenbaum et al. (2020a), and Farboodi et al. (2021), strongly argue in favor of too many meetings, while others, as Rachel (2020), support the view that there could be too much social distancing. The presence and strength of externality are clearly model-dependent and can significantly vary with the details. Thus, we chose to model it in a general reduced form. This allows us to evaluate the effect of its curvature on the optimal intervention. Indeed, in Section 3.3, we show that **conditional on the type** of intervention, the optimal policy is mostly invariant to it.

The function $v(\cdot)$ captures the trade-off between output losses and human lives. This is a controversial component that could play an important role. To isolate its impact, we first compute the optimal paths assuming $v(\cdot) = 0$ so that only the output cost matters. Then, we assume a linear function, and we follow the standard in the literature of calibrating its value to replicate the statistical value of life.

Finally, the value $V(I_t, E_t, N_t)$ would depend on the type of innovation. For instance, with a perfect cure, as in Alvarez et al. (2021), $V(\cdot)$ would only depend on N_t, while if there is no cure but just a vaccine, as in Shimer and Wu (2021), all three arguments are necessary. Nevertheless, as it is evident from problem (P), as long as ϕ is small, the impact of the vaccine is minimal and akin to discounting. We show later that this is indeed the case for empirical relevant values. However, for sufficiently large values of ϕ, it has a non-trivial non-monotonic impact. We argue in Section 5 that it can trigger a "hope for the cure" effect akin to the critical mass: even if $\underline{E} = 0$, the government can switch to a suppression strategy to keep agents alive until the imminent cure arrival and solves the problem once and for all.

It is worth stressing that problem (P) entails solving for the *entire optimal path* of q_t, allowing for a variety of possibilities. This contrasts with other approaches in the literature consisting of comparing alternative scenarios, e.g., no intervention vs. a given fixed quarantine q, or fixing the value of q and comparing the effects of different durations. With our approach, these scenarios would be particular cases.

3 Quantitative implications

This section calibrates the model, explains the calibration strategy and data measurement issues, and uses the calibrated model to assess the optimal quarantine policies.

3.1 Functional forms

We start by specifying the model's functional forms. One of our goals is to analyze how optimal policies change with different degrees of tolerance to output variation. If the welfare function is linear, a complete economy's shutdown could be optimal whenever the cost of doing so is compensated later once the virus is no longer a problem. Instead, when the welfare function has a low intertemporal elasticity of substitution (IES), it is more costly to reduce current economic activity to reap future benefits. We adopt the following function:

$$u(c) = \frac{c^{1-\varepsilon} - 1}{1 - \varepsilon}, \text{ if } \varepsilon \neq 1 \quad \text{and} \quad u(c) = \log(c), \text{ if } \varepsilon = 1.$$

This functional form is mathematically tractable and meaningful in one dimension or another, allowing for a wide range of interpretations. Section 3.3 computes the optimal policy for a wide range of ε, varying it from zero to 4.

The choice of $v(\cdot)$ is contentious. One may think that a quadratic loss function could be appropriate so that the cost rapidly grows with the number of fatalities. However, this function also embodies the unappealing feature that the planner prefers many fatalities if they are sufficiently spread out over time, versus a few concentrated in a short interval. An alternative is a linear function $v(x) = dx$, where d reflects the statistical value of life. In this case, the total number of deaths matters more than when they occur.[16] As Alvarez et al. (2021) and Farboodi et al. (2021), we follow this approach. However, instead of fixing a specific value for d, we consider a wide range, ranging from 0 to 100 years of per capita output. We do this because (1) the literature on the statistical value of life is not yet settled, see Cutler and Summers (2020), and (2) society may value lives beyond their purely statistical value.

We assume $m(\iota, E) = E$ following the current economic literature. Still, because we use the term $\frac{S}{\iota}$, our functional form allows for some congestion, a feature that is absent in Alvarez et al. (2021), Farboodi et al. (2021), and Berger et al. (2020), among others.

Finally, considering the empirical literature on COVID-19, we allow the behavioral response to depend not only on the intensity of the government intervention but also on the

[16] With $\rho = 0$, time becomes irrelevant, and only the total matters. Otherwise, the planner prefers later deaths. We thank an anonymous referee for this clarification.

threat that the virus poses to society. We assume a linear reaction that satisfies

$$b_t = \beta_0 q_t + \beta_1 D_t + \beta_2 \{\exists \text{ COVID-19}\}_t. \tag{8}$$

Equation (8) contains three coefficients. The first and easiest to interpret is β_0, which captures the effectiveness of the government intervention. One could expect a coefficient around 1, i.e., any intervention translates one-to-one into reduced social interactions. Nevertheless, it could be either smaller than 1 when the intervention substitutes other social interactions or larger than 1 when there are complementarities. β_1 captures the continuous response of the population to the virus spread. If there were no deaths, agents would not be concerned and would continue their normal activities. As the prevalence of the virus increases, the population becomes increasingly concerned and interacts less.

Some recent empirical studies, e.g., Durante et al. (2021), have also found that *awareness* of COVID-19 can cause a discontinuous and persistent change in habits. To allow for this possibility, we introduce the indicator $\{\exists \text{ COVID-19}\}_t$, which takes the value 1, if it is known that COVID-19 is present at time t, and zero otherwise. Thus, the last coefficient, β_2, captures the population's sudden change in behavior when they learn about the outbreak: it is a permanent increase in social distancing. As a result, absent government interventions, social distancing would be $\beta_1 D_t + \beta_2$ at time t. This component will play an important role in our no-intervention scenario and in shaping the optimal policy.

3.2 Fitting the virus dynamics

We calibrate the model economy to Italy at a daily frequency. We set $t = 0$ to January 1, 2020. The daily discount rate ρ is set to $0.05/365$ to match an annual interest rate of 5%. Appendix B provides an extensive and detailed explanation of our approach to estimating the model. In general terms, we use clinical data to pin down γ and η. Ferguson et al. (2020) and Chen et al. (2020) provide abundant information about successive illness stages.

To estimate equation (8), we use the movement index developed by Durante et al. (2021), which we complement with a proxy for q_t constructed by Guiso and Terlizzese (2020). Table 1 summarizes the estimated coefficients for equation (8), while the details appear in Appendix B.2. All the coefficients have the expected sign and are highly significant. The awareness coefficient is 0.25: as soon as the population became aware of the virus, they reduced their movements by almost 25%. Then, people took additional precautions as the number of deaths increased. For instance, $1,000$ daily deaths reduced their movements by an extra 11%. This means that movements and social interactions were drastically reduced even before the government took any measures.

Table 1: Social Distancing Equation

Variable	Coefficient	Std. Dev.	t-Stat
β_0: government policy q	1.225	0.0822	14.89
β_1: reaction to deaths	0.00011	0.000034	3.29
β_2: awareness	0.246	0.0312	7.87

The lockdown seems to be very effective: a coefficient of 1.22 implies that the quarantine had an effect beyond the reduction in production, which could be due to complementarities. Policy enforcement could have increased the population's compliance, perhaps feeling that it was their "civic duty," as most of society was enduring the same problems. It could also reflect changes in habits. For instance, shoppers buy more goods on fewer trips.

There are four parameters for which there is ample uncertainty: θ, σ, λ, and $\hat{t}_0^I$: the initial day of the outbreak. We estimate them to match the observed dynamics of COVID-19 in Italy. One approach would be to fit the dynamics of infected cases, but this is problematic: the official reported number of "cases" captures only the individuals who were tested and generated a positive result.[17] Since our model features a one-to-one mapping between infections and fatalities, we target the fatalities path. But also these measures are controversial. Since many fatalities, especially at the peak of the infection, may have been reported as unrelated to COVID-19. To deal with this caveat, following Rinaldi and Paradisi (2020), we use the Italian excess fatalities relative to the five previous years instead. Additional details and the uncovered patterns can be seen in Figure 8 of Appendix B.

To be precise, we target the path of daily fatalities from March 8 to May 31, 2020. We exclude the dates before the first intervention on March 8 (the first reported COVID-19 death was on February 22) because the relevance of COVID-19 for causing excess deaths is almost irrelevant. We are confident, though, that most of the excess deaths in March are COVID-19 related. We choose $\{\hat{t}_0^I, \theta, \sigma, \lambda\}$ to minimize the following loss function:

$$\mathcal{L} = -\sum_{t=t_0}^{T} \left(\frac{D_t^{\text{model}} - D_t^{\text{data}}}{D_t^{\text{data}}} \right)^2, \tag{9}$$

where D_t^{model} is daily fatalities generated by the model and D_t^{data} is the daily death number in the data. Loosely speaking, the initial fatalities between March 1 and March 8 are mainly determined by the outbreak day and λ, the scale parameter in the meetings function. The outcomes after the consecutive government interventions and the population's reaction shed

[17] See Hortaçsu et al. (2021) for an estimation of underreporting in the United States.

Table 2: Parameter Values

Parameter	Symbol	Value	Moment
Preset			
Daily discount rate	ρ	0.05/365	Interest rate
Initial exposed	E_0	2/60mn	Two individuals in population
Critical mass	$\underline{E}$	{0,1/60mn}	Minimum possible number
Hospital capacity	h	0.0067	5,343 ICUs for 60 million
Exposed-to-infected rate	γ	0.143	7-day incubation period
Recovery rate for I	η	0.058	17 days to recovery for I
Jointly calibrated			
Contagion rate	λ	0.400	Fit fatality's path
Recovery rate for E	σ	0.016	Fit fatality's path
Daily death rate if treated	θ	0.06%	Fit fatality's path
Daily death rate if untreated	δ	0.12%	Fit fatality's path
Implied moments			
Basic reproduction number	R_0	2.53	

light on the number of asymptomatic agents, determined by σ and the fatality rate θ.[18]

The variation generated by successive government interventions is instrumental in identifying the parameters (see Figure 3). Three important changes in the intensity of the intervention, plus the initial information on fatalities, provide four moments for the remaining four parameters. The impact of interventions is mainly determined by σ, controlling the speed at which asymptomatic subjects recover. Table 2 shows the calibrated parameter values. We obtain that $\lambda = 0.4$, $\theta = 0.06\%$, $\eta = 0.058$, $\sigma = 0.016$, and the day of the outbreak is January 5, 2020. A few points are worth mentioning. First, our calibrated model's implied initial reproduction factor is $R_0 = 2.53$, which is in line with most of the literature.[19] Second, σ is small compared to η, implying that most E types transition to I and remain asymptomatic for a long time. Because of the E's slow recovery, quarantines can take some time before they have an effect. Finally, the estimated daily death rate implies an average fatality rate of 1.0% for appropriately treated patients.

The comparison between model generated and observed fatalities are plotted in Figure 2. Panel A are the daily flows, while in Panel B the cumulative numbers. Excess deaths are depicted with a red "+" mark and the *reported* COVID-19 deaths with a blue "x" mark. Each mark is one observation. The solid blue line corresponds to the model-generated deaths.

[18] It is key for this argument that most of the symptomatic were quarantined. If there were no asymptomatic ($\sigma = \infty$), the virus would have died out rapidly, while with many asymptomatic ($\sigma = 0$), the infection would have kept spreading quickly.

[19] See Billah et al. (2020) for an extensive review of estimated R_0's worldwide.

Figure 2: Infections and Fatalities: Model vs. Data

Panel A: Daily　　　　Panel B: Cumulative

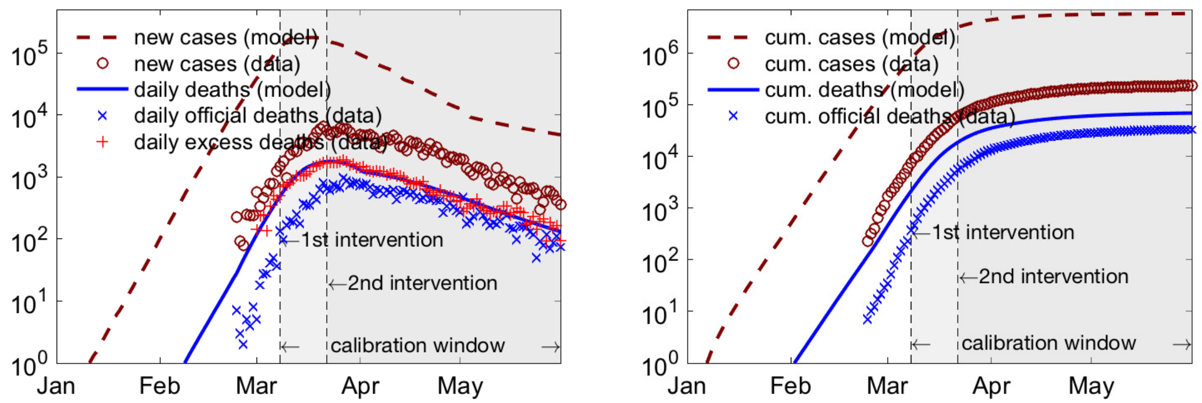

Notes: The vertical axis is in log scale in both panels.

Panel A suggests that the model generated path of daily deaths fits well the measured excess deaths. We also show the infections. The reported infected cases are plotted with burgundy circles, and the model-implied new cases are plotted with a dashed burgundy line. The model-implied cases are substantially above the actual measured cases. To put it in context, at the peak of the spread, the number of officially reported new infections was around $6,550$, while the model suggests that there were about $100,000$ new infections on the same day. The underreporting is also well captured in Panel B with cumulative measures. There are sizeable distances between the model-generated fatalities and cases and their officially reported numbers. For example, on May 3, 2020, the model generated cumulative fatalities is twice the officially reported one. Regarding cases, the discrepancy is more pronounced. The model estimates that by May 3, 2020, more than 2 million people had or have had the virus, while the analogous official number was around $300,000$.

In Appendix C, we present the counterfactual predictions absent policy inventions for completeness. We skip those details to move to our primary focus: optimal policy. Yet, it is worth mentioning two main features of the no-intervention scenario. First, the endogenous behavioral responses imply substantially lower fatalities than those predicted by SIR models abstracting from it. Initial estimates in Walker et al. (2020) projected around $645,000$ fatalities for Italy without intervention, while we obtain $215,000$. Second, the outbreak is so explosive that the economy reached herd immunity by the end of the summer.

3.3 Optimal indiscriminate quarantines without testing

To better understand the consequences of each moving part, we start analyzing the optimal quarantine policy, leaving testing for the next section. We also set $\phi = 0$ so that a vaccine is not expected. We delay its implications for Sections 4 and 5 because, as we then explain, it is either inconsequential shaping q_t or generates relevant strategy switches that deserve detailed explanations. As a reference, if the vaccine is expected to arrive in two years, the arrival rate is $\phi = (1/2)/365 \approx 0.0014$.

The optimal policy is an infinite-dimensional object, stating the intensity of the intervention for all future periods. Let Z_t be a variable encompassing all the information related to the economy's state. We assume a three-parameter flexible functional form of Z_t to simplify the optimization problem. To be precise, q_t is a three-parameter step function such that, for some $\tilde{q} \in [0, 1]$, $\varpi \in \mathbb{R}$, and $\tau \geq 0$, the optimal intervention satisfies:[20]

$$q_t = \begin{cases} \tilde{q} + \varpi \times Z_t, & \text{if } t \leq \tau \\ 0, & \text{if } t > \tau. \end{cases} \tag{10}$$

For instance, a complete shutdown of all economic activities for two weeks would be represented by $\tilde{q} = 1$, $\varpi = 0$, and $\tau = 14$; policies with $\varpi = 0$ would characterize any fixed-intensity intervention. The parameter ϖ captures the time-varying intensity component resembling Z_t's shape. A fitting example can be found in the Italian response to the second wave. The government implemented a "traffic lights" system resembling equation (10). The system considers three-zone colors, white, yellow, and red, with increasing restrictions and minimum transiting times across lights. So, $\tilde{q}$ would reflect the restrictions in the white zone, $\varpi > 0$ the increased restrictions due to a worsening Z_t, and τ the minimum time the region must spend in each light.

Since all variables are linearly related, H_t, E_t, or I_t would capture the shape of Z_t in a similar way, except perhaps, by the nonlinearities introduced by the hospital capacity. Still, although small, the different shapes of the variables could provide important information and help approximate the optimal unrestricted policy better. To address this concern, we estimate the optimal policy assuming the dependency of q_t on H_t, E_t, I_t, and $\Delta_t I_t$. We find that a policy that depends on I_t generates the highest welfare and therefore set $Z_t = I_t$ in equation (10). Nevertheless, the differences among the alternatives are minor.[21]

[20] We started experimenting with a function depending only on time. This time dependency captured the planner's reaction to the state of the economy, such as "health state," e.g., number of cases, number of fatalities, and hospital capacity. Thus, we decided to make the dependency on the state explicit.

[21] This is the same choice as the Italian health authorities for the traffic lights system. It is worth mentioning that we wrote our first draft before this system was implemented.

Figure 3: Optimal Quarantines

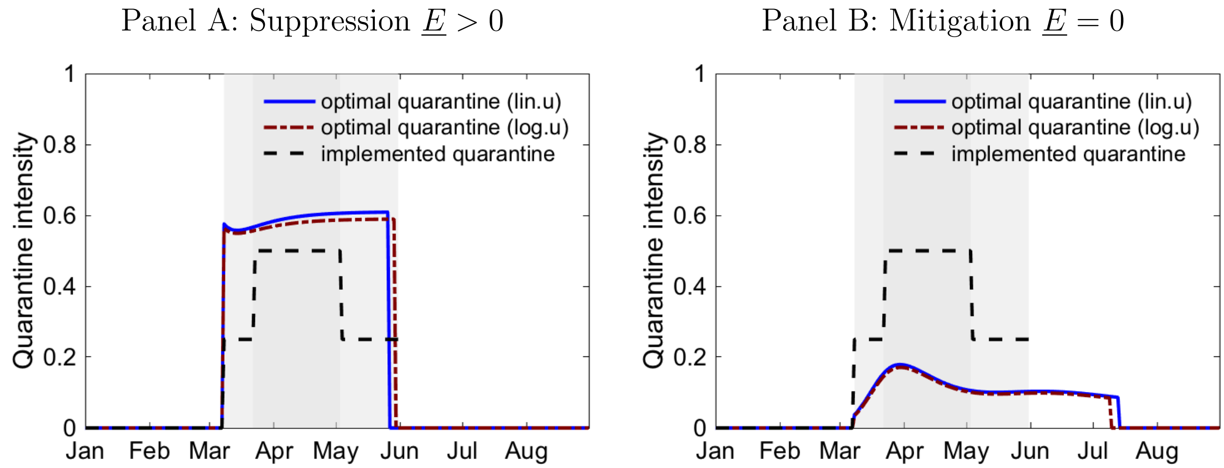

Suppression vs. mitigation

Now, in addition to characterizing the optimal policy without testing, we show the relevance of the critical mass $\underline{E}$. The role of this mass is to undo the "human divisibility assumption," so we set it to the equivalent of one person: $\underline{E} = 1/60\text{million}$. If the exposed number is reduced to less than one individual, the virus disappears. To compare with the observed interventions, we assume that the optimal quarantine is implemented on day 48, corresponding to March 8, 2020: when the Italian lockdown started.

Figure 3 depicts the main takeaway from this section: there are two types of intervention, suppression and mitigation, and **conditional of the type of intervention, the path is mostly invariant to the parameterization of the welfare function**. In Panel A, we plot two representative optimal suppression policies and compare them with the calibrated intensity of the observed lockdown. Suppression consists of a strict lockdown with an approximately constant intensity for a determined time span. Interestingly, the welfare function curvature affects neither the intensity nor the duration. Similarly, in Panel B, we show two representative mitigation policies. Mitigation is less intense than suppression and more responsive to the state of the economy. It starts at moderate levels and slowly decreases over time. Again, the intervention's path is unaffected by the welfare function's curvature.

Figure 3 shows representative optimal paths; under which combination of parameters are they optimal? We find sharp switches between the two policy regimes depending on the underlying model parameters. Since the optimal quarantine q is barely affected by the parameters, conditional on the strategy, we focus instead on total fatalities. Figure 4, plots how total deaths vary with the curvature parameter ε and the statistical value of life d,

Figure 4: Optimal Quarantine Strategies

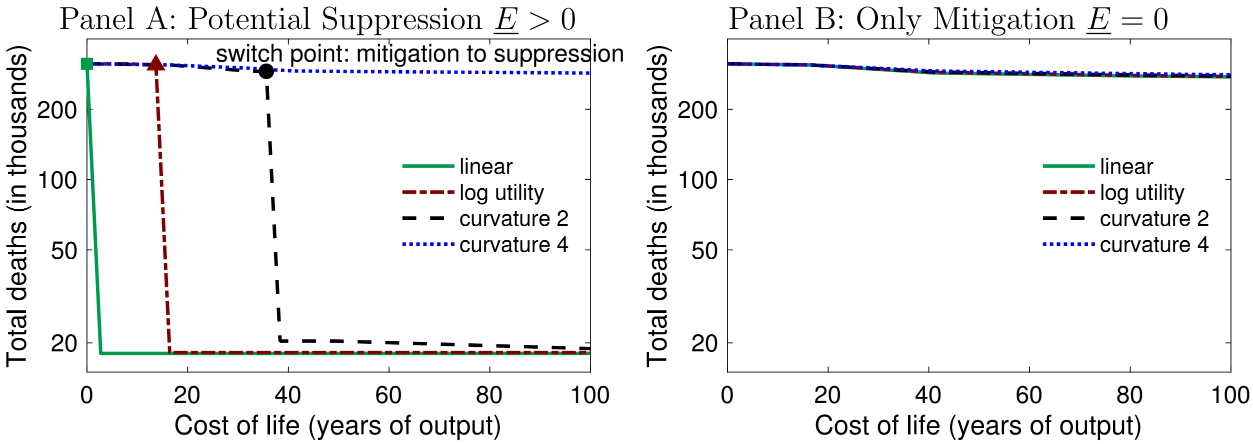

conditional on $\underline{E} > 0$, in Panel A, and $\underline{E} = 0$, in Panel B. The high-mortality outcomes result from mitigation policies, while the low-mortality outcomes result from suppression. Again, conditional on the policy, neither the value of life nor the curvature of the welfare function is important in shaping the dynamics. Different combinations of d and ε determine whether the planner wants to follow a no-intervention, suppression, or mitigation strategy.

Consider Panel A, where $\underline{E} > 0$. When the value of life is $d < 2$, the government does not intervene independently of ε. As d increases, it is optimal to intervene, but depending on the combination of d and ε, the **type of the intervention jumps** to either mitigation or suppression. When the welfare function is linear, the optimal policy quickly jumps to suppression. If it is logarithmic, there is no intervention for $d < 13$ and an intermediate area $13 < d < 16$ where mitigation is optimal, jumping to suppression for $d > 16$. Interestingly, the suppression policy under the logarithmic and linear welfare functions coincide. Similarly, for $\varepsilon = 2$, there is a larger area where mitigation is optimal, $13 < d < 36$. Still, in the suppression area, all policies are similar. In contrast, Panel B shows that when $\underline{E} = 0$, suppression is never optimal, no matter how much value society assigns to lives.

Table 3 reports the outcomes under these strategies. Column (2) reports the optimal **suppression** and column (3) the optimal **mitigation** policy. As a benchmark, column (1) reports the analogous statistics for the scenario without intervention. Note that absent any intervention, the model predicts $313,000$ fatalities, around a third of the projection without behavioral response. The suppression strategy takes the form of a roughly constant intensity around 0.6 for 80 days, which resembles the policy implemented in many countries. This policy effectively reduces the number of symptomatic cases from 41.4% to 2.4% of the population and fatalities from 0.52% to 0.03%.

Table 3: Optimal Quarantine Policies

	No Intervention (1)	Suppression (2)	Mitigation (3)
Quarantine:			
Initial day	-	Mar 8	Mar 8
Duration (days)	-	80	128
Maximum q	-	0.61	0.18
Average q	-	0.59	0.11
Symptomatic rate (per person)	41.4%	2.4%	41.5%
Symptomatic people (million)	24.9	1.5	24.9
Immunity rate (per person)	45.4%	2.7%	45.5%
Immune people (million)	27.2	1.6	27.3
Death rate (per person)	0.52%	0.03%	0.46%
Total fatalities (thousand)	313	18	275
Welfare gain (consumption equiv.)	-	2.4%	0.21%

Notes: Columns (2) and (3) report the welfare gains for linear utility. With log utility, the welfare gains are 2.1% and 0.20%, respectively.

In contrast, the mitigation policy is costly in terms of lives and output. The fatalities are slightly below the no-intervention case, $275,000$, vs. $313,000$. But there is a significant buildup of immune individuals, around 27 million, preventing future waves' arrival.[22]

Taking stock, the intensity and the duration of interventions are mostly determined by the dynamics of the virus. This does not mean that the welfare function is irrelevant: it determines the type of intervention. In Appendix D.1, we show that this is due to the convexity of the welfare function. Loosely speaking, $\underline{E} > 0$ generates a **"hope for the cure"** effect, driving governments to save as many lives as they can until the virus is eradicated. Here, the critical mass generates the effect; we show in Section 5 that the same patterns emerges even when $\underline{E} = 0$, but a "cure" is expected to arrive soon enough. *This could explain why there is large observed heterogeneity in the approaches taken by different countries, but not so much when these different approaches are grouped into the three potential characteristics: no-intervention, mitigation, and suppression.*

[22] See Appendix D for the dynamics of the remaining variables and Figure 12 for the dynamics without government intervention and without behavioral response.

4 Combining testing and quarantines

The main reason giving rise to indiscriminate quarantines is the inability of the policymaker to distinguish exposed from susceptible subjects. If the government knew who the carriers were, it could simply quarantine the exposed and leave everyone else unrestricted. The technology to obtain information through testing is certainly available, but it could be prohibitively expensive to test a vast proportion of the population. Since the immediate output cost of quarantine also appears considerable, it is worth evaluating how much the planner would be willing to spend on testing to reduce the quarantine cost.[23]

Testing has always been considered a fundamental component in controlling pandemics. However, the standard approach by epidemiologists is to rely on contact tracing. This approach is useful at the onset of a pandemic, but it quickly becomes unfeasible when the spread gains speed. For this reason, we focus on random testing. Though it may appear foolish to test blindly without precise direction, it is the simplest mass-testing technology that can be implemented, and we show below that it generates considerable welfare gains.

We divide the exposed population into two groups: the unidentified exposed and the identified exposed, i.e., those designated as *positive* carriers. We maintain the notation E for the unidentified. As before, these individuals are indistinguishable from the susceptible S, as are individuals in the R^u group who recovered without ever exhibiting symptoms. Thus, ex-ante, all individuals in the set $S + E + R^u$ look alike to the planner. To separate them, the government randomly tests subjects in this set. Subjects who test positive are placed in a new group, denoted by E^p, and forced into mandatory quarantine. Since we consider a testing technology that cannot detect antibodies, whether the individual is in S or R^u remains unknown to the tester. To summarize, the total population becomes:

$$N_t = \underbrace{S_t + E_t + R_t^u}_{\text{unidentified}} + \underbrace{E_t^p + I_t + R_t}_{\text{identified}}.$$

To understand the relevance of testing, we analyze the new laws of motion. Suppose a fraction α_t of the unidentified individuals is randomly tested, then $\alpha_t E_t$ individuals are identified as positive carriers. The laws of motion for the two exposed groups are:

$$dE_t = \begin{cases} \left[\lambda \frac{S_t(1-b_t)}{\iota_t} m(\iota_t, E_t(1-b_t)) - (\gamma + \sigma + \alpha_t)E_t\right] dt, & \text{if } E_t \geq \underline{E} \\ -(\gamma + \sigma + \alpha_t) E_t dt, & \text{if } E_t < \underline{E} \end{cases} \tag{11}$$

$$dE_t^p = \alpha_t E_t dt - (\gamma + \sigma) E_t^p dt. \tag{12}$$

[23]We abstract from antibody testing. This technology is available and in use in many countries. However, Obiols-Homs (2020) shows that this testing minimally reduces the spread of the virus.

Equations (11) and (12) capture the *first contribution of testing* to welfare. Recall that only the group E_t can spread the disease, so the more contained this group is, the lower the contagion rate. Comparing (11) with (2), it is evident that testing adds a downward drift α_t to the unidentified exposed group. Absent testing, in equation (2), the E group shrinks only when individuals either become symptomatic, at rate γ, or recover, at rate σ. Now, some individuals are also exiting the E group at rate α_t, when they are identified as contagious and enter the positively exposed group in equation (12).

Unlike the unidentified exposed E, when an identified case E^p recovers, she is known to be immune. Then she joins the recovered group R rather than R^u. While the law of motion for unidentified recovered R^u remains the same, the one for identified recovered R is now modified to:

$$dR_t = (\eta I_t + \sigma E_t^p)\, dt. \tag{13}$$

Equation (13) reveals the *second contribution of testing*. Comparing equation (13) to (6), there is an extra inflow of agents with known immunity at rate σE_t^p. Since the recovered E^p are known to be immune, they rejoin the labor force, reducing the quarantine output costs.

In short, the group of positively tested individuals generates a bulk that reduces the speed of contagion and increases the resources available to cope with the quarantine. This is especially important when the exposed individuals may never be symptomatic. Without testing, they would always be treated as the susceptible population subject to quarantines.

The law of motion for the infected is slightly modified to:

$$dI_t = \left[\gamma \left(E_t + E_t^p\right) - \left(\eta + \Delta_t\right) I_t\right] dt, \tag{14}$$

which differs from the previous law of motion (4) by accounting for the inflow of positively-tested exposed subjects. The law of motion for the total population remains identical to equation (7) since only the symptomatically infected are subject to the risk of death.

The production-feasibility set remains unaltered, with the mass $E_t + S_t + R_t^u$ subject to quarantines q_t are R_t allowed to work. However, some resources must be used to pay the testing costs. Let x be individuals tested at a given instant, then the flow cost is governed by the convex cost function $\Phi(x)$, with $\Phi(0) = 0$, $\Phi'(x) > 0$, and $\Phi''(x) > 0$. Accounting for the cost of tests $x_t = \alpha_t \left(S_t + E_t + R_t^u\right)$, the feasibility constraint becomes:

$$Y_t = (1 - q_t)\left(S_t + E_t + R_t^u\right) + R_t - \Phi\left(\alpha_t \left(S_t + E_t + R_t^u\right)\right).$$

Formally, the planner's problem (P) is modified by choosing the join path of quarantine and testing $\{q_t, \alpha_t : t \geq 0\}$, subject to the new laws of motion and feasibility constraint.

4.1 Testing technology

Since the fundamental dynamics of the economy are not affected by testing, we maintain the previous parametrization (Table 2) and let the quarantine policy q_t take the form specified in equation (10). Two additional objects need to be addressed: the testing policy α_t and the testing cost function $\Phi(\cdot)$. Regarding the testing policy, we assume a structure similar to the quarantine function:

$$\alpha_t = \begin{cases} \tilde{\alpha} + \varpi_\alpha \times I_t, & \text{if } t \leq \tau \\ 0, & \text{if } t > \tau. \end{cases} \tag{15}$$

The testing cost function deserves additional considerations. Ample evidence exists about the unit cost for testing, ranging from \$5 to \$50 per test. These costs are estimated around "normal" levels of testing, but we consider a wide range of possibilities, including testing all the population simultaneously. To capture the fact that the cost could significantly increase due to stretched capacity, we assume that the cost function has linear and exponential components: $\Phi(x) = \varrho_0 x + e^{\frac{\varrho_1}{1-x}} - e^{\varrho_1}$ so that testing the whole population is infinitely costly. If ϱ_1 is sufficiently small, the marginal cost is approximately ϱ_0 for low values of x.

We follow Atkeson et al. (2020) in calibrating the linear component of the cost function and considering potential testing errors. The testing technology embodies two test types. First, the population is *screened* with a fast antigenic test that costs \$5 per unit. This test is not 100% accurate: false positives and negatives will arise. Specifically, 2.79% of subjects who are not infected generate a positive result, while 1.5% of those who are infected generate a negative result. Those who tested positive on the first test take a second *confirmation* (molecular) test that costs \$50. If both tests deliver positive results, the subject is forced into quarantine.[24] We then convert the cost to units of daily output. The GDP per capita in Italy is roughly €34,000 per year, a roughly daily output per capita of \$100. As a result, the linear cost is:

$$\varrho_0 = \frac{\$5}{\$100}\left(S + E + R^u\right) + \frac{\$50}{\$100}\left(0.0279\left(S + R^u\right) + 0.985E\right).$$

Note that the marginal cost depends not only on the number of tests but also on the compo-

[24] We could also consider only the molecular test. This would be a suboptimal implementation of the testing technology, substantially increasing the cost. Unlike Atkeson et al. (2020), we do not consider the possibility that subjects who tested positive do not comply with the quarantine. Enforcement is a relevant problem in many countries, especially the U.S., but not in Italy. Still, there is the test specificity problem; even if it were possible to test 100% of the population, 1.5% would be wrongly categorized as "negative." Hence, some individuals would still be moving freely and spreading the virus. This caveat becomes relevant only if the policy were aiming to screen the entire population.

sition of the testing pool. We specify the exponential parameter to $\varrho_1 = 0.1$. In Appendix F, we consider a steeper exponential cost with $\varrho_1 = 0.5$, reflecting a more constrained testing capacity. With the baseline cost function, the exponential cost grows to the equivalent of the entire daily output if we test 86% of the population and just 49% of the population with the steeper cost function.

4.2 Testing as a substitute for quarantines

As we show in detail in Section 5, there are many combinations of parameters for which the potential arrival of a vaccine does not affect the optimal policy. There are, however, some cases in which it does so. For this reason, we present results mostly for the case in which the vaccine is not expected to arrive, and we stress when it matters. We follow the approach by Alvarez et al. (2021) and assume that the vaccine is also a perfect cure. Thus, when it is discovered, all the people still alive remain so. We experimented with the approach of just a vaccine, with no cure for the already infected, but this only makes it even less likely that the vaccine can impact the optimal policy.

In Table 4, we present the main results. We compare the optimal combination of quarantine and testing policies in columns (3), (4), and (6). To contrast with Section 3.3, we reproduce the (only) indiscriminate quarantine results in columns (1) and (2). In column (5), we present a combination of parameters for which the vaccine's expected arrival in two years has a large impact. Unlike the previous section, the curvature parameter is relevant in shaping the optimal policy, but only when following a suppression strategy. Conditional on mitigation, the optimal policy is still invariant to the curvature. For this reason, we have included two alternative computations for the optimal suppression with testing in columns (3) and (4) while presenting only a representative result for mitigation.

Testing is used intensively in all scenarios; average testing during interventions ranges from 8% to 90% of the unidentified population per day. How testing is distributed over time depends on the planner's aversion to output variation. When output variation is not a concern, the interventions are short-lived but very intense. In contrast, when smoothing is a concern, the interventions are long-lasting and less intense. It is also clear from Table 4 that the intense testing is compensating a major reduction in quarantines, either on their intensity or duration. Thus, testing is substituting indiscriminate quarantines. [25]

The significant substitution between quarantines and testing can be seen all across the board. When mitigation is intended, the intensity of the quarantine drops down to zero.

[25] Alvarez et al. (2021) find that random testing would never be optimal. In their computation, they bound testing: it is not possible to test "*at a speed such that the entire population will be tested in a year.*" This bound discards all the paths that we find optimal.

Table 4: Optimal Quarantine and Testing Policies

	Quarantine Only		Quarantine & Testing			
	Suppress.	Mitiga.	Suppression			Mitigation
			Lin. Util	Log Util	Log Util	
	(1)	(2)	(3)	(4)	(5)	(6)
Critical mass $\underline{E}$	1 person	0	1 person	1 person	0	0
Vax. arrival ϕ, annual	0	0	0	0	0.5	0
Intervention:						
Initial day	Mar 8	Mar 8	Mar 8	Mar 8	Mar 8	Mar 8
Duration (days)	80	128	1	18	1016	362
Quarantine:						
Maximum q	0.61	0.18	0.58	0.17	0	0
Average q	0.59	0.11	0.58	0.09	0	0
Testing:						
Maximum α	-	-	0.90	0.63	0.52	0.26
Average α	-	-	0.90	0.63	0.13	0.08
Total cost (% of GDP)	-	-	0.4%	1.1%	6.9%	1.3%
Sym. rate (per person)	2.4%	41.5%	2.4%	3.0%	3.8%	41.4%
Sym. people (million)	1.5	24.9	1.5	1.8	2.3	24.8
Asym. rate (per person)	-	-	0.13%	0.12%	0.2%	0.8%
Asym. people (thousand)	-	-	72	99	124	480
Immune rate (per person)	2.7%	45.5%	2.7%	3.3%	4.2%	45.4%
Immune people (million)	1.6	27.3	1.6	1.9	2.5	27.2
Death rate (per person)	0.03%	0.46%	0.03%	0.038%	0.05%	0.44%
Tot. fatalities (thousand)	18	275	18	23	29	266
Welfare gain (c. equiv.)	2.4%	0.21%	3.0%	2.9%	2.1%	0.46%

Notes: Columns (1) and (2) report the welfare gains for linear utility. With log utility, the welfare gains are 2.1% and 0.20%, respectively.

When instead suppression is the goal, either the quarantine's intensity or duration are substantially reduced, and testing is significantly scaled up. In any case, rather than enacting indiscriminate and inefficient quarantines, the planner prefers to use massive testing to identify the virus' carriers and restrict them, and only them. As expected, the cost is not negligible, with a minimum amounting to 0.4% of GDP, depending on the parameters configuration. Nevertheless, it is definitively smaller than reducing daily production by 60% for over 80 days (around 13% of annual GDP) under the optimal suppression strategy, or reducing daily production by 11% for over 128 days (around 4% of annual GDP) under the optimal mitigation strategy.

These numbers are far larger than the observed testing intensities by many countries,

Figure 5: Optimal Quarantine and Testing

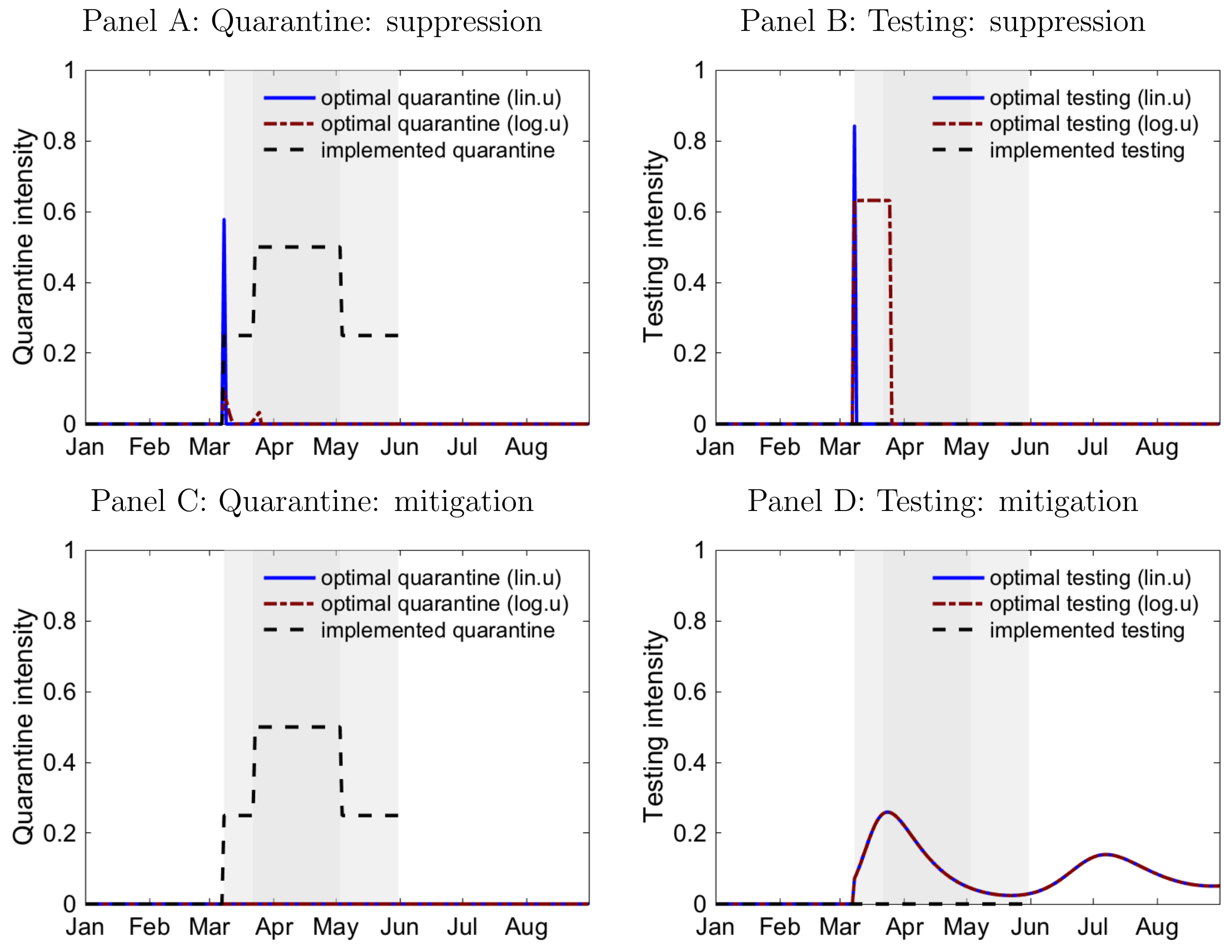

except for China. In all cases, sizeable welfare gains accrue compared to no testing scenarios. With the logarithmic function, the consumption-equivalent welfare gain is 2.9%, compared to only 0.21% without testing. The important takeaway from this result is that *testing is a substitute for, not a complement of, quarantines*, reducing either the intensity or the duration of quarantines. Looking at overall fatalities and infections, it is evident that the outcomes are very similar to those in which testing is not allowed. **The main difference lies in the path for output, which generates larger welfare gains**.

The quarantine and testing paths are depicted in Figure 5. The top panels, A and B, correspond to the suppression strategies in columns (3) and (4) of Table 4. The bottom panels, C and D, depict the mitigation strategy referred in column (6) of Table 4. In all cases, the solid blue line represents the case with a linear welfare function, while the dashed brown line represents the logarithmic welfare function. Two things are fairly evident from this figure. 1) When mitigation is the intended strategy, independently of the welfare function's curvature, indiscriminate quarantines are eliminated and substituted by massive testing. 2) If suppression is possible, there is still a large substitution of lockdowns for testing, but the extent of it depends on the aversion to output variation. For instance, if the welfare function is linear (blue line in panels A and B), both the quarantine and testing spike for one day, testing the entire unidentified population.[26] To summarize, when testing is not possible, a mitigation strategy is characterized by a long-lasting, low-intensity quarantine, while the suppression strategy is consists in shorted duration by higher intensity intervention. *When testing is possible, all the intensity is loaded to testing rather than quarantines.*

The implied output, number of unidentified asymptomatic individuals, and the identified asymptomatic individuals by each optimal testing strategy can be found, respectively, in Panels A, B, and C of Figures 14 and 15 in Appendix E. Panel C presents an additional measure that exists only with testing: asymptomatic individuals previously identified as positive and then recovered. Because now it is known that they are immune, they can work. This measure becomes particularly relevant with long quarantines. After three months, it amounts to almost 2% of the labor force.

It is worth noting that the mitigation strategy, due to its smoothing features, implies a long-lasting prevalence of the virus. The second wave of testing observed in Panel D of Figure 5 mirrors the future increase in cases due to yet unachieved herd immunity. Note also that the Italian implemented quarantine reaches suppression levels, but without the necessary testing to eliminate the virus once and for all, it falls short of an optimal suppression and log of optimal mitigation.

Finally, in column (5) can be seen that the expected arrival of a vaccine (when it matters) may have a significant impact reshaping the optimal policy. To understand why, it is useful to compare it with columns (4) and (6), where all parameters are the same except ϕ. When $\phi = 0$, since $\underline{E} = 0$, the government follows a long-lasting intervention consisting of very intense testing without quarantines. As with all standard mitigation policies, these policies generate sizeable fatalities. When $\phi > 0$, the government scales up significantly the testing efforts while still not using quarantines. The effort significantly reduces the number of fatalities and the immunity rate by an order of magnitude. Indeed, comparing it with column (4), this strategy, rather than pursuing mitigation, is following suppression despite $\underline{E} = 0$: eliminating the virus without resorting to herd immunity. **Suppression is achieved by completely substituting quarantines with testing**.

[26] We want to emphasize that the percentage is with respect to the unidentified $S + E + R^u$, not with respect to the entire population, so that the number of tests is continuously decreasing over time.

5 Medical innovations: hoping for the cure

At the moment of the outbreak, there was a major debate about the vaccine timeline. Based on previous experiences, many specialists argued that it was too optimistic to think about it as a possibility before three years. Others argue that something could be done with significant investment in just a year. Ex-post, the vaccine arrived in the lower end of the estimates.[27] How do these ex-ante expectations affect the optimal policy? If different health authorities have different expectations, would they react differently?

Given the previous results, one may wonder whether the illustrative example we showed Section 4 represents a general pattern or just a particular case that we cherry-picked. We argue in this section that for intermediate values of ϕ, the expectation of a potential future vaccine creates an *option value* akin to the critical mass: a hope for the cure. In turn, the hope that the disease can be eradicated adds value to those policies that preserve lives during the vaccine waiting time.

However, the option value is non-monotonic in ϕ. Thus, it may either generate significant changes on the optimal policy or being innocuous. The non-monotonicity of the option value is fairly intuitive. On the one hand, suppose that $\phi \to \infty$ so that the cure is expected to arrive immediately, then the illness does to pose a threat and the option value is zero. On the other hand, suppose that $\phi = 0$ so that the cure will never arrive, then even though it would be valuable, since the vaccine would not arrive its value is again zero. It is for high uncertainty configurations (intermediate ϕ) that this effect kicks in. Then, how does it translate into policies? We show here that when there is no critical mass ($\underline{E} = 0$), there is a non-monotonic relationship between the intervention's intensity and ϕ; the optimal policy may jump from mitigation to suppression. Instead, when the critical mass is already present ($\underline{E} > 0$) the effect is monotonic. Since a "hope" with option value is already present, the intervention's intensity is mostly unaffected by ϕ, as long as it is not too large.

To understand these arguments consider Figure 6, which plots the optimal lockdown policies when testing is not available for different values of ϕ. These values range from $\phi = 0$, so that the vaccine never arrives as in Section 3.3, to an annual rate of 365, in which case the vaccine is expected to arrive the next day. In all cases, the welfare function is linear. Panel A displays the average intensity of the quarantine, Panel B the duration, and Panel C resumes this information depicting the number of fatalities **conditional** on the cure never arriving. In each panel, the continuous blue line represents an economy with a critical mass equivalent to 1 person, while the dashed brown line an economy without critical mass.

[27] Arrival must be interpreted carefully. One thing is the discovery and another the implementation. Many countries are still literally vaccine-less.

Figure 6: Only quarantines: hope for the cure effect

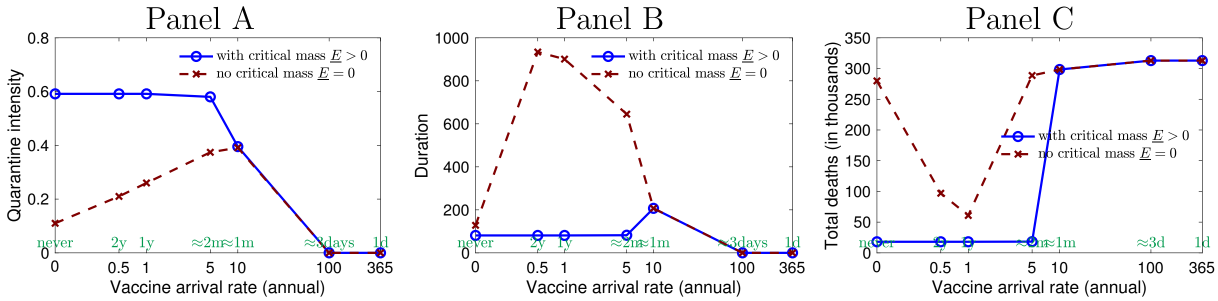

Note: The horizontal axis is in log scale in the three panels.

When the vaccine is expected to arrive almost immediately, it is never optimal to intervene irrespective of the critical mass. Since the vaccine is on the way and the virus takes time to reproduce, it is better to let it spread and avoid the output cost. Even though expected, the vaccine may never arrive, in which case the number of fatalities is large (Panel C). Technically speaking, the vaccine (or cure), has zero option value. As the expected time of arrival increases (ϕ decreases), the option value becomes positive, then it is optimal to intervene preserving lives until the cure solves the problem once and for all. Intuitively, for sufficiently short expected arrival times, the intervention is not expected to last long, implying a small output cost. Hence, the interventions with and without the critical mass coincide: the policy is intended to suppress the virus.

Note, however, that as ϕ decreases the paths conditional on $\underline{E}$ diverge. If $\underline{E} > 0$, the policies quickly converge (for an expected arrival larger than two months) to the analogous with $\phi = 0$ as in Section 3.3, while if $\underline{E} = 0$ the policies follow a non-monotonic pattern. In this last case, the option value is sufficiently large that it creates an effect similar to a critical mass, deepening the intervention and saving more life. In one case, the hope is in the hands of the government (belief in critical mass), while in the other depends on innovators (belief in science). This effect eventually fades out; for very low values of ϕ, the cure is no longer expected, so the option value is again zero, converging to the results in Section 3.3.[28]

The switch from mitigation to suppression is more striking when testing is possible, as depicted in Figure 7, illustrating the optimal quarantine and testing policies conditional on $\underline{E} = 0$. The panels' interpretation is similar to Figure 6, but there is an extra panel depicting the testing intensity. Moreover, the continuous blue line now corresponds to a linear welfare function, while the brown dashed line is a logarithmic one. The non-monotonic effect of the option value is very clear. Also, the switch to the suppression strategy for intermediate values

[28]Few papers have taken a deeper look at these vaccine implications. See Garriga et al. (2020) and van Wijnbergen (2021).

Figure 7: Testing and quarantines: hope for the cure effect: $\underline{E} = 0$

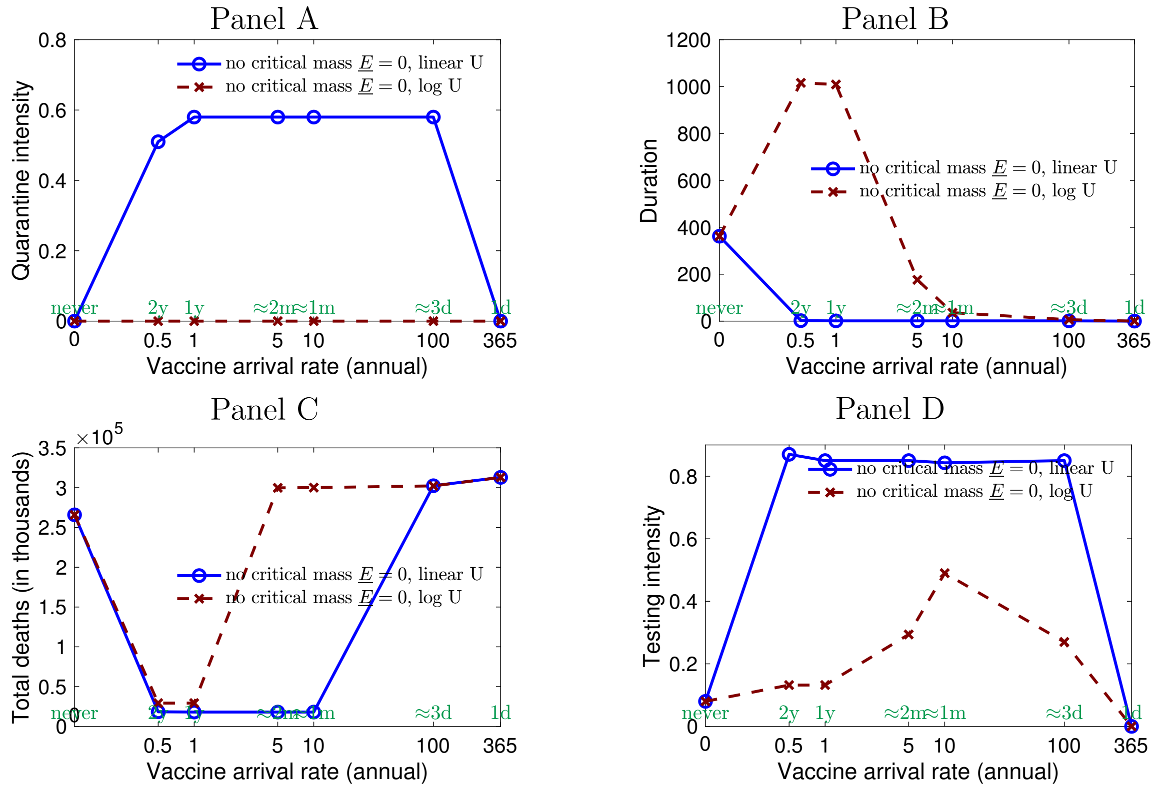

Note: The horizontal axis is in log scale in all panels. Deaths are computed conditional on the vaccine not arriving and are the equivalent to excess deaths in data rather than the officially reported COVID fatalities. See Figure 16 for the case with $\underline{E} > 0$

of ϕ. Recall that since $\underline{E} = 0$ suppression is never optimal when $\phi = 0$. As in Section 4, the suppression strategy manifests in massive testing rather than a harsh lockdown. Indeed, when the welfare function is logarithmic, the quarantine is minimal, loading the weight of the intervention on intense testing. In Appendix G, Figure 16, we show the analogous patterns when $\underline{E} > 0$. There it is clear that, because of the hope-for-the-cure effect is always present, the optimal policies are monotone in ϕ: there is always the option to eliminate the virus.

5.1 Heterogeneous responses and beliefs about the cure

One important takeaway from the last two sections is that minor parameter variations can generate drastic changes in the optimal implemented policies. It is widely believed that some countries choose "soft" (or non at all) lockdowns because the authorities, or the society, put a higher weight on the economic activity rather than lives. Our analysis shows that this intuition is in some way correct. However, this tradeoff can rationalize only a few cases where countries with a very low valuation of life would decide not to intervene, while the other would most likely follow some form of mitigation strategy. If the expectation is that the only definitive solution lies in herd immunity, then the government needs to make sure that hospitals are not saturated. After this is taken care of, the necessary fatalities to obtain herd immunity are already determined and are unavoidable. Moreover, since quarantine paths are determined by the virus' dynamics, one should observe a large cluster of countries following similar mitigation strategies.

Instead, there has been a wider range of observed heterogeneity in the responses. There are some emblematic examples like China and New Zealand with very strict lockdowns, compared with the light measures imposed in Sweden and many U.S. states. Although it is possible to argue in favor of no-intervention appealing to the lives-output trade-off alone, strict lockdowns can only be rationalized if hope for a cure exists. This hope can manifest in two alternative ways. The government may believe that it could eliminate the virus by stopping all human interactions for a determined time span. Once the last person infected is either recovered or dead, the virus is gone for good. This interpretation could suit well in countries that can control their borders or have a legal system more amicable with restrictions to free movements.

The hope can also stem from believing that either a vaccine or effective treatment would soon arrive and solve the problem. Unlike the mitigation case, now the country can achieve immunity without paying a stiff price in human lives: deaths are avoidable. Thus, the government can have incentives to implement strict measures to wait out for the vaccine to arrive. This process could save many lives that, after the cure's arrival, would be permanently safe. One country that could fit this profile would be Norway. Of course, for this strategy to succeed and make sense, the expected arrival time must not be too distant into the future.

As a result, if one groups countries according to their initial beliefs about the possibility of reaching the critical mass and their beliefs about the expected waiting time for the vaccine, each group would be very different from the other. Moreover, once quarantines are combined with testing, the heterogeneity of suppression strategies grows considerably (see Table 4). Unlike mitigation strategies, the suppression strategy paths do depend on the trade-off between lives and output when testing is possible. However, one should expect that as the virus spreads across locations, the belief that it could be eliminated should fade away. Similarly, as the outcomes of medical research efforts become public, the beliefs about the vaccine should become more homogeneous.

6 Conclusions

In this paper, we have extended the standard epidemiological SIR model allowing for asymptomatic subjects to be tested, considering the trade-off with output losses. Our results are twofold. First, we show that the specification of the welfare function has little impact on the optimal lockdown path policy. Slight alterations in the trade-off between fatalities and output may generate jumps in the government's strategies. However, it does not change the optimal path once the strategy is chosen. The jumps are due to the policymaker's believes about the possibility that the virus could be eliminated or that a cure would eventually arrive and do the job. This finding could explain the widely diverse policies followed by different countries. Some, such as China, New Zealand, and Australia, chose strict lockdowns. In contrast, others, such as the United States, Sweden, and Brazil, have relied on mitigation strategies, with higher implied costs in terms of lives but lower output costs.

Second, we find that strong reliance on lockdowns alone is inefficient. Random mass testing may appear untenable and extremely costly initially, but it is a feasible, vastly superior alternative to the even more expensive indiscriminate restrictions on economic activity. Pandemics are not new to humanity, and the risk of future ones remains. When they happen, they can be devastating. An important takeaway from our paper is that if the friction is the lack of information, policies should be directed to overcoming it. Testing is superior to lockdowns, not just a complementary policy but a tool that can eliminate the output costs.

References

Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2021, December). Optimal targeted lockdowns in a multigroup sir model. *American Economic Review: Insights 3*(4), 487–502.

Alvarez, F., D. Argente, and F. Lippi (2021, September). A simple planning problem for covid-19 lock-down, testing, and tracing. *American Economic Review: Insights 3*(3), 367–82.

Assenza, T., F. Collard, M. Dupaigne, P. Feve, C. Hellwig, S. Kankanamge, and N. Werquin (2020). The hammer and the dance: Equilibrium and optimal policy during a pandemic crisis. Cepr discussion papers n 14731, CEPR.

Atkeson, A. (2020, March). What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. Working Paper 26867, National Bureau of Economic Research.

Atkeson, A., M. C. Droste, M. Mina, and J. H. Stock (2020, October). Economic benefits of covid-19 screening tests. Working Paper 28031, National Bureau of Economic Research.

Bar-On, Y., T. Baron, O. Cornfeld, R. Milo, and E. Yashiv (2021, March). COVID-19: Erroneous Modelling and Its Policy Implications. IZA Discussion Papers 14202, Institute of Labor Economics (IZA).

Barrot, J.-N., B. Grassi, and J. Sauvagnat (2021, May). Sectoral effects of social distancing. *AEA Papers and Proceedings 111*, 277–81.

Berger, D., K. Herkenhoff, C. Huang, and S. Mongey (2020). Testing and reopening in an seir model. *Review of Economic Dynamics*.

Billah, A., M. Miah, and N. Khan (2020, Nov). Reproductive number of coronavirus: A systematic review and meta-analysis based on global level evidence. *PLoS One 15*(11).

Brotherhood, L., T. Cavalcanti, D. Da Mata, and C. Santos (2020, August). Slums and Pandemics. CEPR Discussion Papers 15131, C.E.P.R. Discussion Papers.

Brotherhood, L., P. Kircher, C. Santos, and M. Tertilt (2020, May). An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies. CEPR Discussion Papers 14695, C.E.P.R. Discussion Papers.

Chari, V. V., R. Kirpalani, and C. Phelan (2021, October). The Hammer and the Scalpel: On the Economics of Indiscriminate versus Targeted Isolation Policies during Pandemics. *Review of Economic Dynamics 42*, 1–14.

Chen, J., T. Qi, L. Liu, Y. Shi, T. Zhu, and H. Lu (2020, May). Clinical progression of patients with covid-19 in shanghai, china. *Journal of Infection 80*, E1–E6.

Cutler, D. M. and L. H. Summers (2020, 10). The COVID-19 Pandemic and the 16 Trillion Virus. *JAMA 324*(15), 1495–1496.

Dewatripont, M., M. Goldman, E. Muraille, and J.-P. Platteau (2020). Rapid identification of workers immune to covid-19 and virus-free: A priority to restart the economy. Discussion paper, Universit Libre de Bruxelles.

Durante, R., L. Guiso, and G. Gulino (2021). Asocial capital: Civic culture and social distancing during covid-19. *Journal of Public Economics 194*, 104342.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020a). The macroeconomics of epidemics. Working Paper 26882, National Bureau of Economic Research.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020b, May). The macroeconomics of testing and quarantining. Working Paper 27104, National Bureau of Economic Research.

Farboodi, M., G. Jarosch, and R. Shimer (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory 196*(C).

Favero, C. A., A. Ichino, and A. Rustichini (2020, April). Restarting the economy while saving lives under Covid-19. CEPR Discussion Papers 14664, C.E.P.R. Discussion Papers.

Ferguson, N. M., D. Laydon, and G. Nedjati-Gilani (2020). Impact of non-pharmaceutical interventions (npis) to reduce covid-19 mortality and healthcare demand. On behalf of the imperial college covid-19 response team, Imperial College of London.

Garriga, C., R. Manuelli, and S. Sanghi (2020, May). Optimal Management of an Epidemic: Lockdown, Vaccine and Value of Life. Working Papers 2020-031, Human Capital and Economic Opportunity Working Group.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Ríos-Rull (2020, April). Health versus wealth: On the distributional effects of controlling a pandemic. Working Paper 27046, National Bureau of Economic Research.

Greenwood, J., P. Kircher, C. Santos, and M. Tertilt (2019, July). An Equilibrium Model of the African HIV/AIDS Epidemic. *Econometrica 87*(4), 1081–1113.

Guimarães, L. (2021). Antibody tests: They are more important than we thought. *Journal of Mathematical Economics 93*(C).

Guiso, L. and D. Terlizzese (2020). Sfruttare la chiusura per riaprire prima possibile. *Il Foglio* (March 24th).

Hansen, C. H., D. Michlmayr, S. Madeleine, P. K. Mølbak, and S. Ethelberg (2021, March). Assessment of protection against reinfection with sars-cov-2 among 4 million pcr-tested individuals in denmark in 2020: a population-level observational study. *The Lancet 397*, 1204–1212.

Hortaçsu, A., J. Liu, and T. Schwieg (2021). Estimating the fraction of unreported infections in epidemics with a known epicenter: An application to covid-19. *Journal of Econometrics 220*(1), 106–129. Pandemic Econometrics.

Kaplan, G., B. Moll, and G. L. Violante (2020, September). The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the u.s. Working Paper 27794, National Bureau of Economic Research.

Kermack, W. O. and A. G. McKendrick (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character 115*(772), 700–721.

Moser, C. A. and P. Yared (2020, April). Pandemic Lockdown: The Role of Government Commitment. NBER Working Papers 27062, National Bureau of Economic Research, Inc.

Obiols-Homs, F. (2020). Precaution, social distancing and tests in a model of epidemic disease. Barcelona GSE Working Paper 173, Barcelona GSE.

Pollinger, S. (2020, May). Optimal Case Detection and Social Distancing Policies to Suppress COVID-19. TSE Working Papers 20-1109, Toulouse School of Economics (TSE).

Rachel, L. (2020, December). An analytical model of covid-19 lockdowns. Discussion Papers CFM-DP2020-29, Centen for Macroeconomics.

Rinaldi, G. and M. Paradisi (2020, April). An empirical estimate of the infection fatality rate of covid-19 from the first italian outbreak. *medRxiv*.

Shimer, R. and L. Wu (2021, March). Diffusion on a sorted network. Working paper, University of Chicago.

Smith, L. E., H. W. W. Potts, R. Amlôt, N. T. Fear, S. Michie, and G. J. Rubin (2021). Adherence to the test, trace, and isolate system in the uk: results from 37 nationally representative surveys. *BMJ 372*.

van Wijnbergen, S. (2021, May). Lockdowns as options. Tinbergen Institute Discussion Papers 21-037/IV, Tinbergen Institute.

Walker, P. G., C. Whittaker, et al. (2020). The global impact of covid-19 and strategies for mitigation and suppression. On behalf of the imperial college covid-19 response team, Imperial College of London.

Wilder-Smith, A., C. J. Chiew, and V. J. Lee (2020). Can we contain the covid-19 outbreak with the same measures as for sars? *The Lancet Infectious Diseases 20*(5), e102–e107.

Appendix

A Model moments

For an infected patient, if treated, the density function of dying after s units of time is:

$$f^t(s) = \theta e^{-(\eta+\theta)s}.$$

The death rate is:

$$\int_0^\infty f^t(s)\, ds = \int_0^\infty \theta e^{-(\eta+\theta)t} ds = \frac{\theta}{\eta+\theta}.$$

The density function of recovering after s units of time is

$$g^t(s) = \eta e^{-(\eta+\theta)s}.$$

The average recovery duration is:

$$\int_0^\infty g^t(s)\, sds = \int_0^\infty \eta e^{-(\eta+\theta)s} sds = \frac{\eta}{(\eta+\theta)^2}.$$

Similarly, if untreated, the death rate is $\frac{\delta}{\eta+\delta}$. The average recovery duration is $\frac{\eta}{(\eta+\delta)^2}$.

B Parameters estimation

We first calibrate the parameters that have a direct medical interpretation and for which there is clinical information. The transition rate γ, from exposed to infected, directly maps to the incubation period, which is 6.5 days according to Ferguson et al. (2020). Thus, we set $1/\gamma = 7$. The recovery rate of symptomatic subjects, η, is related to the recovery time. This moment, however, falls into a wide range. Ferguson et al. (2020) state that it takes nine days on average for a subject to recover. In contrast, the data on active cases in Italy implies that it takes an average of 48 days. The latter number appears exaggerated, probably reflecting delays in the administrative process to "officially" declare a subject as recovered.[29] The experience for the European outbreak shows that nine days seems to be on the low end; many studies suggest somewhere around 14 days (see Chen et al. (2020)). For this reason, we target an average recovery time of 17 days.

[29] An individual is declared recovered after two consecutive tests with negative results. Thus, delays in testing could lead to the delayed resolution of the active cases. Also, some regions take longer than others to process the patients' paperwork.

Specifically, in our model the recovery time is coupled with the death rate and depends on whether the infected patients are treated or not. When treated, a patient recovers in $\frac{\eta}{(\eta+\theta)^2}$ days, and a fraction $\frac{\theta}{\eta+\theta}$ of patients die. When left untreated, a patient recovers in $\frac{\eta}{(\eta+\delta)^2}$ days, and a fraction $\frac{\delta}{\eta+\delta}$ of patients die. Details for the computation are included in Appendix A. Knowing the daily death rate θ, we can then back out the recovery rate η to target a recovery time of 17 days.

There are four parameters for which there is ample uncertainty: θ, σ, λ, and $\hat{t}_0^I$: the initial day of the outbreak. The daily death rate is also needed to pin down the recovery time. σ determines how long an asymptomatic agent can be contagious and therefore is important shaping the dynamics. Finally, λ, determines the reproduction factor of the virus R_0, which has been the subject of considerable debate. One may think that σ is a fundamental property of the virus, in the sense that it should be the same across countries. However, both θ and λ can be specific to each country. For instance, the age structure of the population can affect average death rates, while the nature of the social interactions would determine λ.

To address these issues, we estimate $\{\hat{t}_0^I, \theta, \sigma, \lambda\}$ to match the observed dynamics of COVID-19 in Italy. One approach would be to fit the dynamics of infected cases, but this is problematic: the official reported number of "cases" captures only the individuals who were tested and generated a positive result. There are many reasons to believe that this measure underrepresents the true cases. First, those asymptomatic are rarely tested and therefore not recorded. Second, as is well known, test kits were scarce during the onset of the pandemic, which forced the authorities to test only subjects who were likely to be infected or vulnerable. Thus, many mildly symptomatic individuals were left untested.

Our model features a one-to-one mapping between infections and fatalities, so we target the fatalities path. But also these measures are controversial. Since many fatalities, especially at the peak of the infection, may have been reported as unrelated to COVID-19, this statistic could be underestimated. To deal with this caveat, and following Rinaldi and Paradisi (2020), we compute the excess daily fatalities in Italy relative to the previous years. In Panel A of Figure 8, we show the daily excess number of deaths and the daily official deaths due to COVID-19. It is evident that the steep increase in excess deaths at the peak of the first wave, in March 2020, is not reflected in the official death number. This pattern remains for most of the sample, albeit with a reduced gap after the initial peak. For this reason, we are confident that the excess death number provides a more accurate measurement.

To be precise, we target the path of daily fatalities from March 8 to May 31, 2020. We exclude the dates before the first intervention on March 8 (the first reported COVID-19 death was on February 22) because the relevance of COVID-19 for causing excess deaths is lower. We are confident, though, that most of the excess deaths in March are COVID-19

Figure 8: Data on Fatality and Movements

Panel A: Fatalities   Panel B: Movements

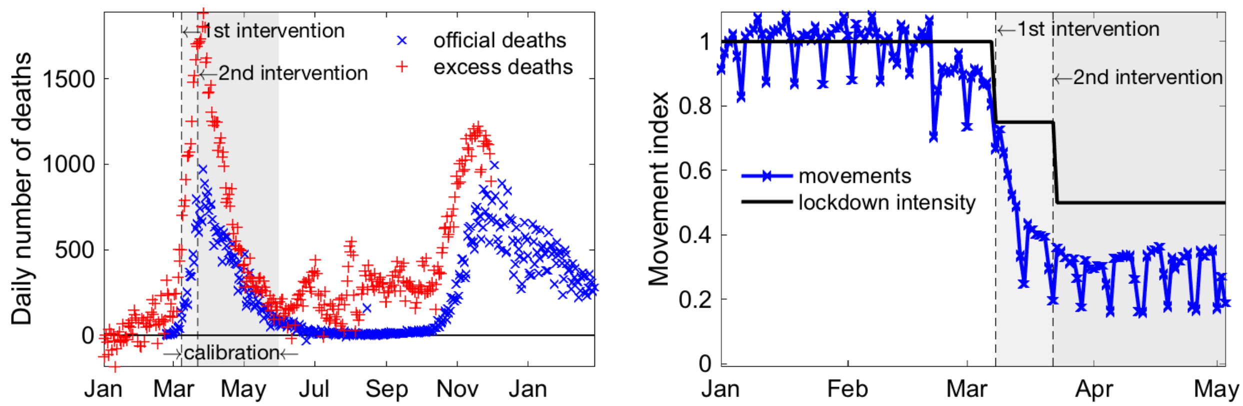

Notes: In Panel A, daily excess deaths are computed as the number of deaths relative to the average of the same day in the previous five years, based on the information released by *Instituto Nazionale di Statistica*: found at https://www.istat.it/it/archivio/240401. The plotted data series ends on November 30, 2020. Since the number of deaths in January 2020 is relatively smaller than the previous years, we normalized the series such that the average for January is zero. In Panel B, the movement series is normalized such that the average movement before the arrival of COVID-19 is 1. For further details, see Durante et al. (2021).

related. This leads to the choice of $\{\hat{t}_0^I, \theta, \sigma, \lambda\}$ to minimize the loss function in equation (9).

To implement this calibration strategy, we need three additional pieces of information. First, for a given number of initially recorded fatalities, many combinations of outbreak date and initially exposed mass are consistent with the observed initial deaths. To avoid this ambiguity, we set the initial infection to two persons, i.e., $E_0 = 2/60\text{million}$.[30]

Second, as we specify in equation (8), the extent of social distancing depends on government interventions and the population's behavioral response. To measure government intervention, we appeal to some estimates of the effect of lockdown policies on economic activity using information from Guiso and Terlizzese (2020). They estimate that the initial intervention, on March 8, affected 16% of sectors and the second intervention, on March 22, reached 40% of sectors. We adjust these values upward, because these estimations do not consider the effect of school closings. As stated by Barrot et al. (2021), the fact that workers remained home to take care of their children had an important impact on GDP. Thus, we impose that the initial intervention is $q = 1/4$ and the second is $q = 1/2$.[31]

We measure social distancing using the movement index by Durante et al. (2021), based

[30] The Italian population is around 60 million people. The initial reports in Italy estimated that there were two independent outbreaks, the "Lombardia" and "Veneto" clusters.

[31] We could directly use the ex-post measured GDP. However, the resulting output would be endogenous to the behavioral responses. Our measured q, which uses the ex-ante shares of each affected sector, instrumentalizes the intensity of the intervention.

on cellphone information. The observed mobility pattern is plotted in Panel B of Figure 8 with the blue curve. For comparison, the black line represents the implemented lockdown measures q_t. A couple of features are worth noting. First, it is evident that individuals started reducing their movements before the restrictions were imposed (also documented by Farboodi et al. (2021) in the U.S.). The initial endogenous reaction prior to March 8 is measured by the distance between the blue and black lines. Second, restrictions further reduced mobility. The initial mild lockdown introduced on March 8 is accompanied by a sizable reduction in movements; the more intense lockdown introduced on March 22 led to a further decline. This approach delivers the results in Appendix B.2 summarized in Table 1.

Finally, we assume that hospital capacity started to bind on February 24, 2020 and continued to do so until March 31, 2020. As hospital capacity binds, for the untreated patients, we assume that the daily death rate is double of those treated, i.e., $\delta = 2 \times \theta$. In Appendix B.1, we describe in detail the rationales leading to this assumption. This implies that initially the number of fatalities was larger because some patients were left untreated.

Our setting does not explicitly distinguish between patients who require critical care versus those who don't. Thus, we scale the observed hospital capacity to compensate for this feature. At the time of the outbreak, there were $5,343$ beds in intensive care units (ICU) for a population of 60 million.[32] Since only 1.32% of infections need critical care, the country can treat no more than $5,343/0.0132$ infected individuals at a time.[33] Thus, the country is prepared to treat only $(5,343/0.0132)/60\text{million} = 0.67\%$ of the population. To capture the observed increase in capacity, we assume a quadratic function for the ratio H_t/I_t such that $H_t < I_t$ before February 24 and after March 31, with the series reaching a minimum on March 14 (greatest excess demand for ICU beds). This choice makes sure that the excess capacity pastes smoothly in the corners and that the dates exactly pin down the parameters of the quadratic function. The dates are chosen to match the sharp increase in deaths during March and the subsequent fall in April (see Panel A of Figure 8). Given I_t, we can recover the implicit series for hospital capacity H_t.

B.1 Estimating fatality rates

There is much controversy about the "true" value of the fatality rate, especially when all the available data is too raw to provide a concrete answer. Most studies tend to state that on average 1% of the infected die. However, this value could significantly change with the

[32] At the time of the outbreak, this was announced by the Italian prime minister and confirmed later by data releases. Since October 2020, the Italian government has been posting daily figures about ICU capacity, but for the time frame of our calibration, this information is not readily available.

[33] This calculation is based on facts provided by Ferguson et al. (2020): 4.4% of those infected with COVID-19 become hospitalized and 30% of hospitalized patients would require critical care.

Table 5: Fatality Rates: South Korea and Italy

Classification		South Korea						Italy					
		Cases		Fatal cases			Weight	Cases		Deaths		Lethality	Weight
		Number	(%)	Number	(%)	Rate (%)	age group	Number	(%)	Number	(%)	(%)	age group
All		9,137	100	126	100	1.38	1	35,731	100	3,047	100	8.5	1
Age	Above 80	406	4.4	55	43.65	13.55	0.0342	5,352	15	1,243	50.2	23.2	0.0717
	70–79	611	6.7	39	30.95	6.38	0.0672	7,121	19.9	1,090	35.8	15.3	0.0988
	60–69	1154	12.6	20	15.87	1.73	0.1198	6,337	17.7	312	10.2	4.9	0.1216
	50–59	1724	18.9	10	7.94	0.58	0.1648	6,834	19.1	83	2.7	1.2	0.1549
	40–49	1246	13.6	1	0.79	0.08	0.1626	4,396	12.3	25	0.8	0.6	0.1531
	30–39	943	10.3	1	0.79	0.11	0.1405	2,525	7.1	9	0.3	0.4	0.1172
	20–29	2473	27.1	0	0	0	0.1327	1,374	3.8	0	0	0	0.1027
	10–19	475	5.2	0	0	0	0.0954	270	0.8	0	0	0	0.0956
	0–9	105	1.2	0	0	0	0.0828	205	0.6	0	0	0	0.0843

demographic structure of the population. In particular, the fatality rate appears to sharply increase with age and the lack of proper treatment. Given our focus on the outbreak in Italy, both factors are first-order issues for our estimation.

The study by Ferguson et al. (2020) estimates that the average fatality rate in Wuhan is around 0.99%. They also state that around 4.4% of the infected subjects require hospitalization and that 30% of the hospitalized cases require critical care; even with proper critical care, a patient dies with 50% chance. If we assume that without critical care the subject dies with certainty, it implies that the fatality rate for the untreated is twice the analogous for the treated. This indicates that $\delta \approx 2 \times \theta$, providing the first support for our calibration.

Another approach to determine the difference between θ and δ is to compare fatality rates in a country loose healthcare capacity with another that was constrained. A candidate for the former is South Korea, while Italy is a clear case of the latter. Table 5 presents the detailed rates by age for the two countries. To compute the average, we add across the age-specific rates weighted by the corresponding weights of the age groups in the population. We obtain that the average rate is 1.22% for South Korea and significantly higher, at 4.09% for Italy. But, how much of the gap is due to older demographics in Italy and how much is due to the overwhelmed Italian healthcare system? To address this question, we made an intermediate calculation: we recompute the average death rate for South Korea using the population weights of Italy, which delivers 1.92%. We interpret the difference $1.92\% - 1.22\% = 0.7\%$ as the pure age composition effect. This number is by itself substantial and informative about the significant risk that COVID-19 presents for an "old" country like Italy.

Still, the case fatality rate in Italy is around 4% and, if our adjustment is correct, the residual two percentage points are not explained by the age composition. If we were to attribute the residual difference to the shortage of proper medical care, we would obtain again that $\delta \approx 2 \times \theta$. Since these two independent sources deliver consistent estimates, we

Table 6: Estimated Coefficients for Social Distancing

Variable	Coefficient	Std. Dev.	t-Stat
$\hat{\beta}_0$: government policy q	1.225	0.0822	14.90
$\hat{\beta}_1$: reaction to deaths	0.00011	0.000034	3.29
$\hat{\beta}_2$: awareness	0.246	0.0312	7.87
Intercept	0.141	0.0305	4.62
Monday	-0.149	0.0394	-3.79
Tuesday	-0.167	0.0394	-4.24
Wednesday	-0.162	0.0394	-4.11
Thursday	-0.172	0.0394	-4.37
Friday	-0.195	0.0388	-5.02
Saturday	-0.115	0.0388	-2.96

calibrate our model with $\delta = 2 \times \theta$.

B.2 Empirical model for social distancing

We estimate equation (8) using the following empirical specification:

$$b_t = \hat{\alpha} + \hat{\beta}_0 q_t + \hat{\beta}_1 D_t^{data} + \hat{\beta}_2 Awareness_t + \sum_{i=1}^{6} \hat{\beta}_{2+i}\ day_t^i, \tag{16}$$

where b_t is social distancing measured as 1−movement index, D_t^{data} is the observed number of fatalities on day t, $Awareness_t$ is an indicator function that takes the value 1 starting on February 20, 2020, when it became widely known that the virus was in the country, and day_t^i captures the day effect from Monday to Saturday. As discussed in Section 2, equation (16) contains three important coefficients. First, β_2 captures the break in behavior due to the outbreak of the virus. It is a permanent and constant increase in social distancing.[34] Second, the parameter β_1 captures the additional reaction of the population as the virus spreads. Thus, absent any government intervention, mobility will reduce by $\beta_1 D_t + \beta_2$ at day t. Third, β_0 captures the effectiveness of the intervention. With full compliance, β_0 should be around 1. But it could be $\beta_0 < 1$, when there is substitution of other social interactions, or bigger than 1 when there are complementarities.

[34] Following Durante et al. (2021), we set the initial "awareness" day as February 20, 2020.

C Simulation without intervention

We can estimate the evolution of the illness, and its economic impact, absent any government intervention, i.e., with $q_t = 0$ for all t. Figure 9 shows the main variables dynamics. In the top panel, we show the endogenous social distancing measures. Even without intervention, social distancing is still reduced with q_t^s falling sharply. This is the blue curve vs. the observed red $*$ data points. Mobility reduces by more than 50%, even without a quarantine. This happens because of voluntary social distancing: without an indiscriminate quarantine, more people are infected, which in turn strengthens the individual reactions ($\beta_1 > 0$).

In the second panel, we plot the proportion of infected and exposed subjects and the hospital capacity at each day. The economy starts with an initial mass of 2/60million of exposed individuals and 0 infected. Initially, the exposed move around and engage in economic activities without necessarily knowing that they are carriers. Soon after, some "confirmed" infected start to arise, but still those numbers are very small, and definitively smaller than the number of exposed individuals. In this period, the growth rate of the infection is high, around 100% per day, but the quantities do not seem alarming due to the still small number of affected individuals. After 45 days, the number of infected is about the same as the number of exposed. At this point, if a policymaker takes a picture of the situation, she can see only the type I individuals, but the number of carriers is $2 \times I$.

Nevertheless, the number of infected cases is still small, although growing over time. The initial slope is steep, with the growth in the number of total cases on an explosive path. The situation deteriorates after around 40 days when the hospital capacity is reached. The fatality rate that was low at the beginning starts to rise due to the infected people who are either untreated or badly treated (third panel). After 90 days, the number of infected is at its maximum, with around 5% of the population symptomatically infected, and 4% infected but not showing symptoms yet.

At this point, the growth rate of the infected starts to decrease. The main reason for this is that the number of susceptible people reached a point, such that the reproduction rate of the virus $\lambda(S/N)$ is lower than its death rate, given by $\eta + \Delta_t$. After that, the virus starts to die by itself. The cost of lives lost is high—without intervention, 0.36% of the population dies, with an analogous effect on total production and consumption.

These simulations yield two important takeaways. First, the virus needs unaffected individuals to reproduce. As the infection spreads, the number of susceptible subjects decreases. More meetings start to happen between exposed and already immune individuals. It is true that still some new subjects become infected, but every period fewer individuals are becoming infected than the people who are either recovering or dying. In addition, the population

Figure 9: Potential Path: No-Intervention Scenario

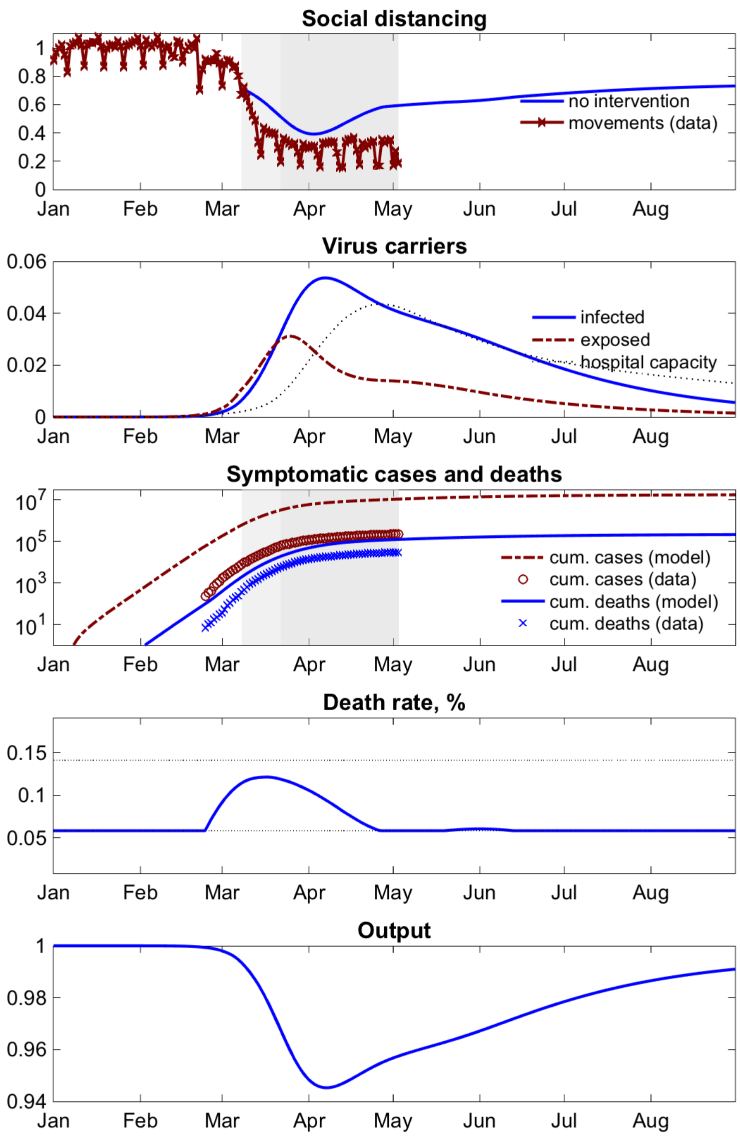

changes its behavior. As the prevalence explodes, agents start to adjust their behavior which significantly slows down the reproduction of the virus. This additional uncoordinated social reaction is key to determining the necessity of further public intervention.

D Supplementary figures: quarantines without testing

In Figures 10 and 11, we plot the effect of the optimal interventions described in Table 3 over output (Panel A), infected cases (Panel B), and fatalities (Panel C). All the patterns are intuitive. The optimal suppression policy generates a large drop in output but also quickly reduces the number of cases and fatalities. In contrast, mitigation has a smaller impact on production but accumulates more infected cases and generates more fatalities; it also achieves the goal of delaying the number of cases and fatalities until the hospital capacity is increased.

Figure 10: Intervention and Effects: Suppression

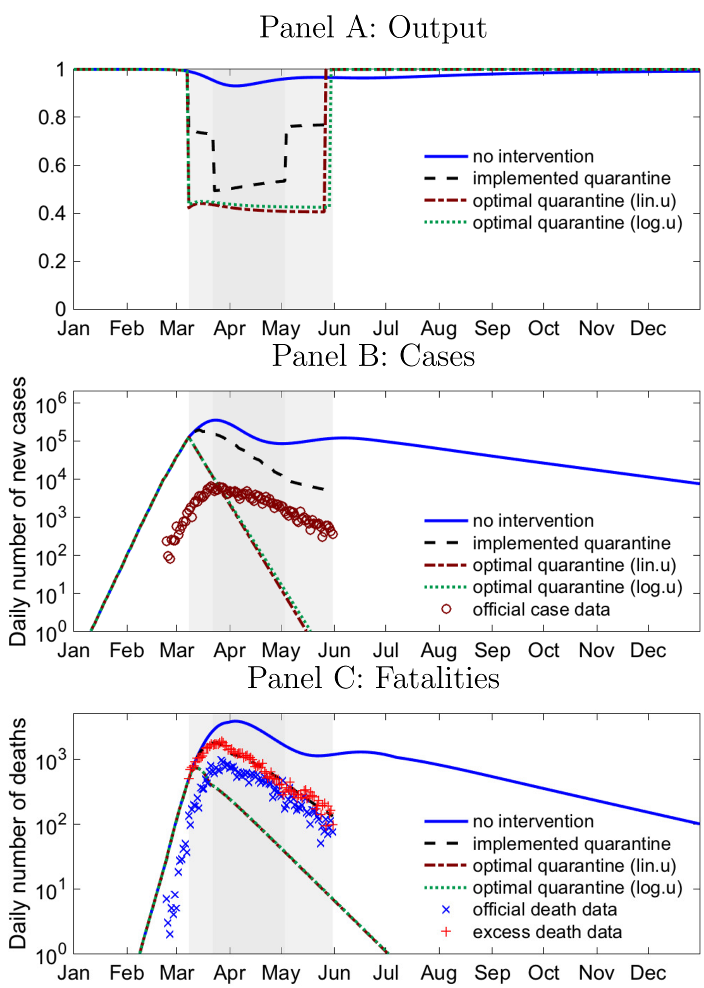

Figure 11: Intervention and Effects: Mitigation

Panel A: Output

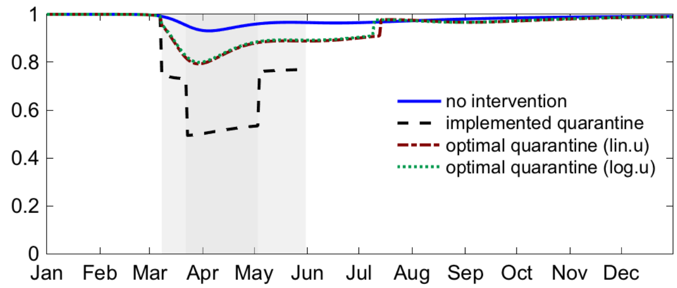

Panel B: Cases

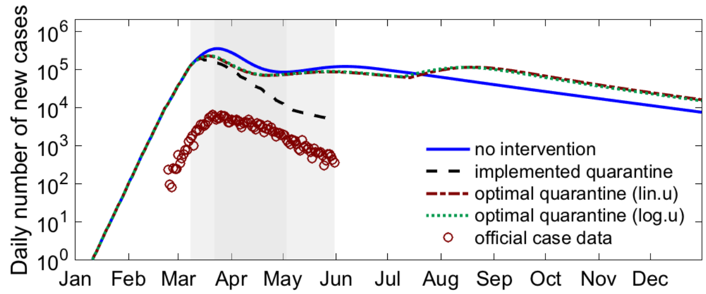

Panel c): Fatalities

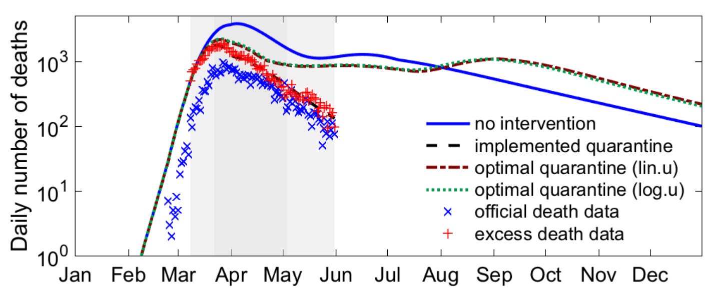

Figure 12: No intervention and no behavioral response

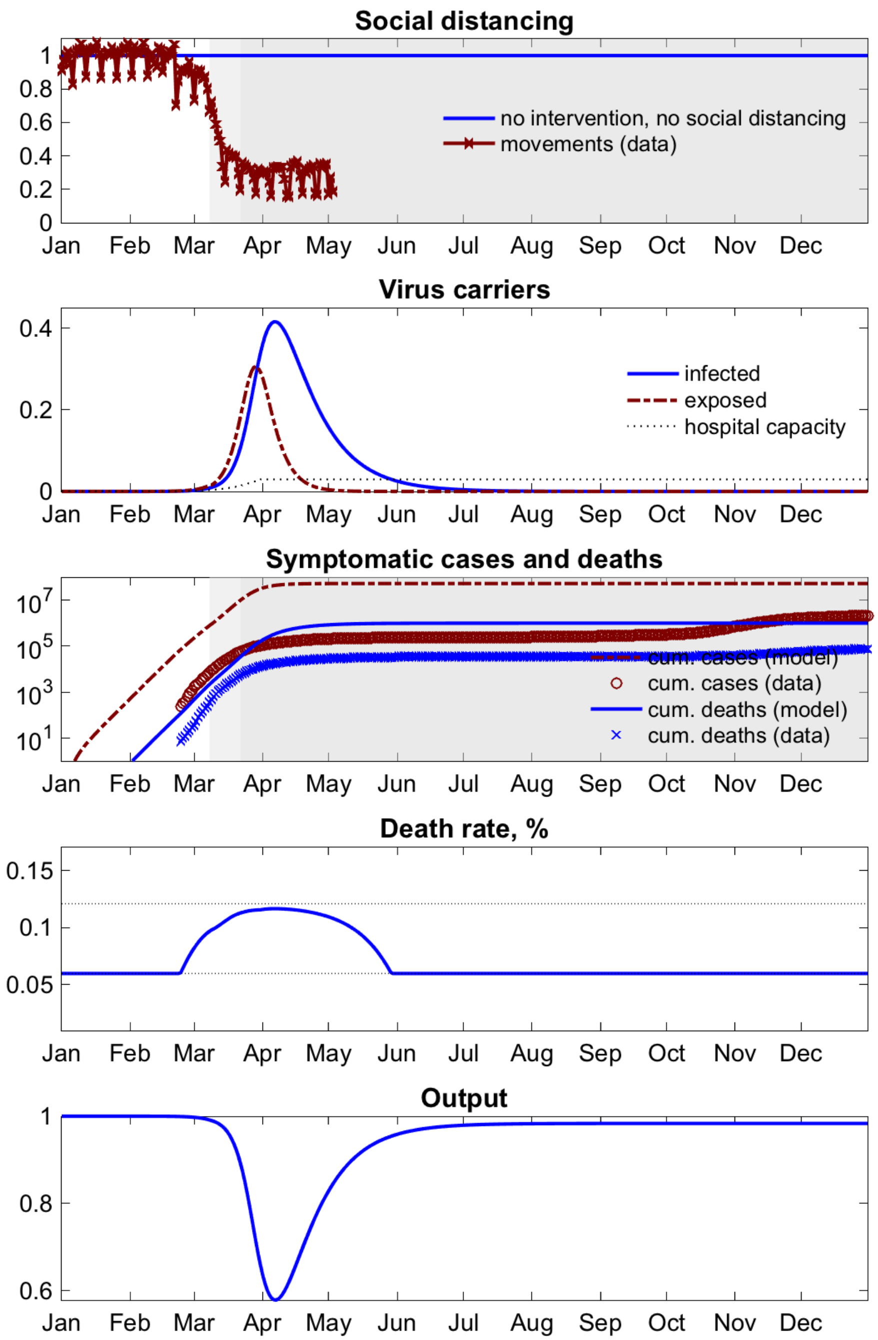

D.1 A convex welfare function

Figure 13 illustrates the lack of concavity in the welfare function depicting a scenario in which the optimal policy is about to switch from mitigation to suppression (the black dot in Figure 4). It plots in three dimensions the welfare outcomes for different combinations of average quarantine intensity and its duration. The welfare function has two peaks, corresponding to the two potential strategies. The mitigation peak corresponds to an intervention with low intensity (11%) and a longer duration (128 days) than the suppression peak, with higher average intensity (59%) and a shorter duration (80 days). The resulting welfare gains are very similar. Changes in the parameters described in Figure 4 change the relative level of the peaks, but not their positions in the figure. That is what generates sudden jumps in strategies with small effects in the implied dynamics. Figure 13 also stresses the difficulty of finding the optimal policy in this kind of setting. Since the problem is not concave, we can not rely on first-order conditions, which can lead to a local maximum. Instead, a grid-search algorithm is guaranteed not to fail, and it delivers the correct global maximum. Moreover, since the exact parametric values lead to either one strategy or the other, it is impossible to state which strategy is better without a good assessment of the cost of lives lost. For this reason, in what follows, we present both types of optimal policies.

Figure 13: Switch between Mitigation and Suppression: Two Peaks in Welfare

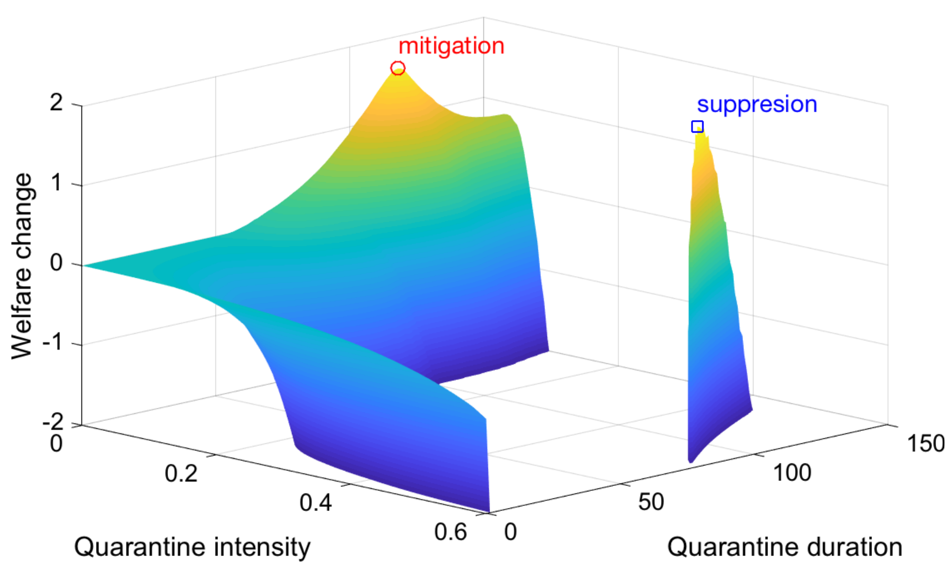

Notes: The changes in welfare are under $\varepsilon = 2$ and $\underline{E} = 1/60$million. The cost of life is 36 years of per capita output, which corresponds to the switching point of the optimal policy from mitigation to suppression as depicted in Panel A of Figure 4.

E Supplementary figures with testing

Figure 14: Intervention and Effects: Suppression with Testing

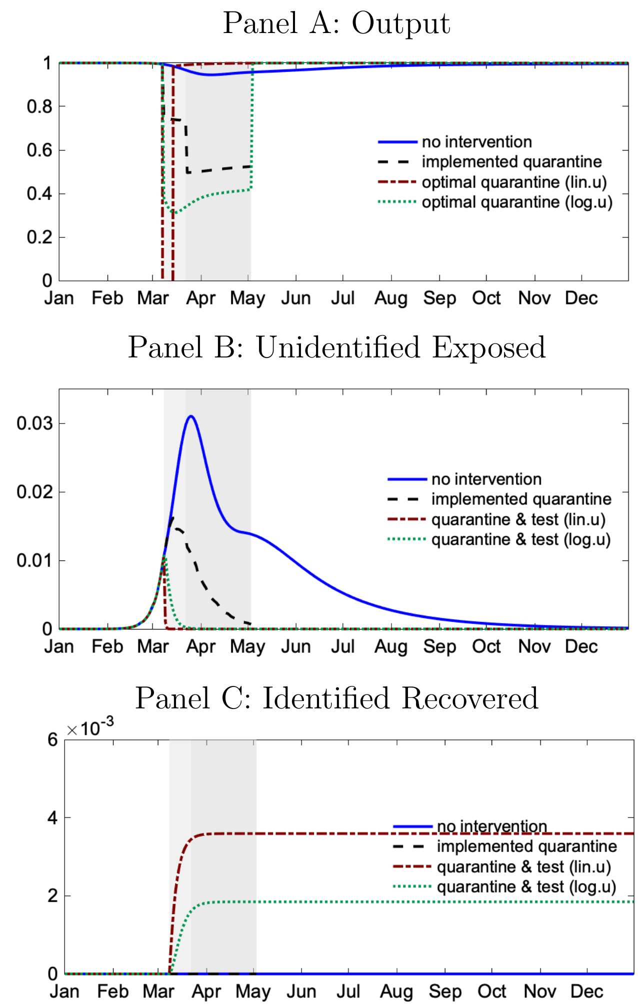

Figure 15: Intervention and Effects: Mitigation with Testing

Panel A: Output

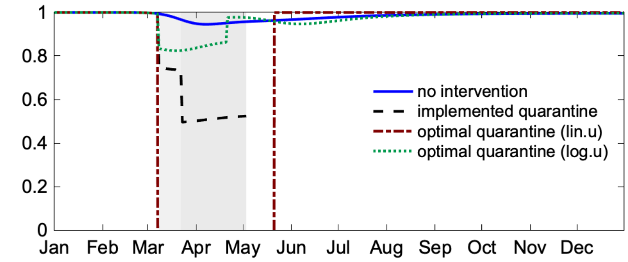

Panel B: Unidentified Exposed

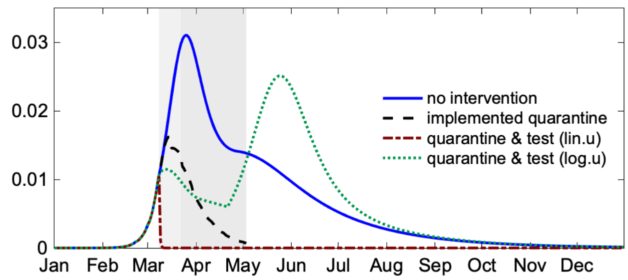

Panel C: Identified Recovered

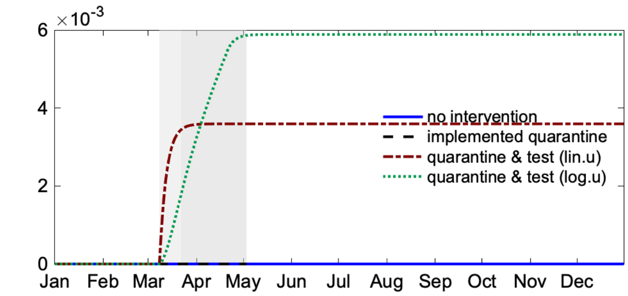

F Higher cost of testing

Table 7: Optimal Quarantine and Testing Policies: High Cost

	Quarantine Only		Quarantine & Testing		
	Suppression	Mitigation	Suppression		Mitigation
			Lin. Util	Log. Util	
	(1)	(2)	(3)	(4)	(5)
Intervention:					
Initial day	Mar 8	Mar 8	Mar 8	Mar 8	Mar 8
Duration (days)	80	128	15	46	321
Quarantine:					
Maximum q	0.61	0.18	0.58	0.22	0
Average q	0.59	0.11	0.56	0.21	0
Testing:					
Maximum α	-	-	0.44	0.28	0.13
Average α	-	-	0.44	0.27	0.04
Total cost (% of GDP)	-	-	3.2%	4.2%	3%
Symptomatic rate (per person)	2.4%	41.5%	2.4%	3.3%	40.9%
Symptomatic people (million)	1.5	24.9	1.5	2.0	24.6
Asymptomatic rate (per person)	-	-	0.11%	0.14%	0.81%
Asymptomatic people (thousand)	-	-	63	90	484
Immunity rate (per person)	2.7%	45.5%	2.7%	3.6%	44.9%
Immune people (million)	1.6	27.3	1.6	2.2	26.9
Death rate (per person)	0.03%	0.46%	0.03%	0.04%	0.46%
Total fatalities (thousand)	18	275	18	25	274
Welfare gain (consumption equiv.)	2.4%	0.21%	2.7%	2.5%	0.25%

Notes: Columns (1) and (2) report the welfare gains for linear utility. With log utility, the welfare gains are 2.1% and 0.20% respectively.

G Vaccines and critical mass together: $\underline{E} = 1/60mn$

This figure is analogous to Figure 7 but with $\underline{E} > 0$ equivalent to 1 individual in the population. The main take away is that the intensity of the intervention is monotone in ϕ.

Figure 16: Testing and quarantines: hope for the cure effect

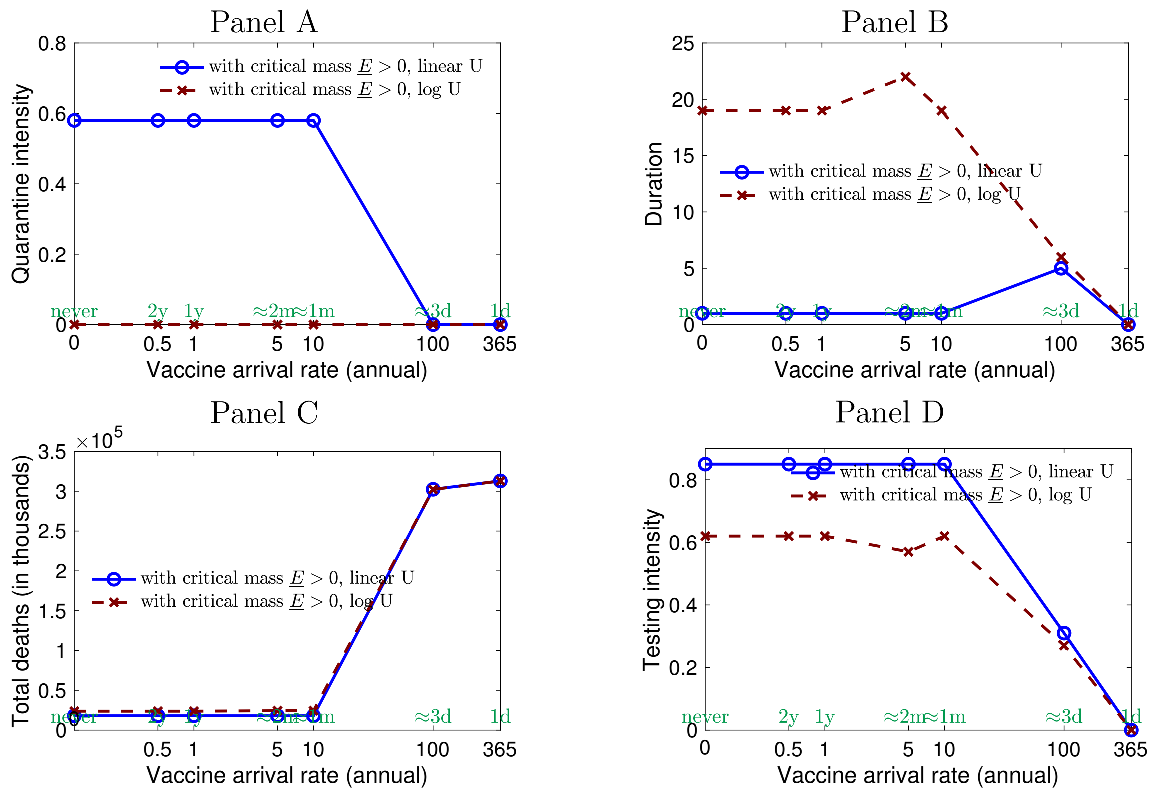

Note: The horizontal axis is in log scale in all panels.