

Behavioral epidemiology: An economic model to evaluate optimal policy in the midst of a pandemic^{*}

Shomak Chakrabarti

Ilia Krasikov

Rohit Lamba

January 2022

Abstract

This paper combines a canonical epidemiology model of disease dynamics with government policy of lockdown and testing, and agents' decision to social distance in order to avoid getting infected. The model is calibrated with data on deaths and testing outcomes in the United States. It is shown that an intermediate but prolonged lockdown is socially optimal when both mortality and GDP are taken into account. This is because the government wants the economy to keep producing some output and the slack in reducing infection is picked up by social distancing agents. Social distancing best responds to the optimal government policy to keep the effective reproductive number at one and avoid multiple waves through the pandemic. Calibration shows testing to have been effective, but it could have been even more instrumental if it had been aggressively pursued from the beginning of the pandemic. Not having any lockdown or shutting down social distancing would have had extreme consequences. Greater centralized control on social activities would have mitigated further the spread of the pandemic.

"Coronavirus is Germany's greatest challenge since World War II, says Angela Merkel", [Deutsche Welle](#) [2020]

"Covid-19 restrictions not affecting social distancing, says ONS: UK statistics agency says level of people still meeting likely to lead to increased hospitalisations and deaths." [Financial Times](#) [2020]

"Vietnam abandons zero-Covid strategy after record drop in GDP: Warnings that lockdowns were crippling businesses heaped pressure on Communist government." [Financial Times](#) [2021b]

I Introduction

Three instruments have been salient in the global response to the Covid-19: government policy of lockdown and testing, and people's decision to practice social distancing. The objective for the collective has largely been the minimization of direct mortality while ensuring a steady pace of the economy, and the objective of the individual has been lowering the chance of getting infected while ensuring some normalcy in life.

While the academic literature in epidemiology is primarily concerned with understanding the evolution of the disease and how to control it, economists are attempting to understand the co-evolution of mortality and economic

^{*}Chakrabarti: shomak.chakrabarti@manchester.ac.uk, University of Manchester; Krasikov: krasikovis.main@gmail.com, Higher School of Economics Moscow; Lamba: rlamba@psu.edu, Pennsylvania State University. We are grateful to David Argente, Shoumitro Chatterjee, Krishna Dasaratha, Elisa Giannone, Callum Jones, Kei Hirano, Shouyong Shi, Shamim Sinnar and Jakub Steiner for helpful comments and suggestions.

output using some subset of the aforementioned three instruments as choice variables subject to some capacity constraint. In that vein, this paper builds a model of disease dynamics where government tries to manage mortality and the economy through lockdown and testing policies, and agents social distance balancing the chance of getting infected with maintaining some social interactions.

While lockdown is good at reducing the spread of infection, it is often accompanied with severe economic costs which may further jeopardize livelihoods even after the pandemic recedes. Social distancing as a behavioral response to the spread of the disease influences the government's choice of severity of the lockdown. And, testing spreads more information all around as to who should be participating in social and economic activities.

We start from an epidemiological framework with five possible states—susceptible, infected, hospitalized, recovered and dead. On it we build a model of economic and social interactions where people are matched and the infection spreads. Economic activities produce output, and can be restricted by the government's lockdown policy. Social activities provide utility to the agents, and are controlled endogenously by their social distancing decisions. Agents suffer disutility from getting infected and infecting others. Both types of interactions— economic and social— can be mitigated by tracing and testing and those who are found to be infected are forced to stay at home. The pandemic can end with the arrival of a vaccine at a random time distributed between one and two years since the inception with a mean of about a year and a half. Finally, anticipating the behavioral response of the agents, the government chooses the optimal testing and lockdown policies, by maximizing a weighted sum of total output and mortality. Each agent in turn takes the government's lockdown and testing policy and other agents's social distancing decisions as given and choose their best response of social distancing.¹

The main contribution of the paper is to model and analyze the three instruments together and explore their pairwise substitutability. This constitutes a technical challenge because we are solving for (i) a planning problem of a forward looking government, (ii) a strategic problem between forward looking agents, and (iii) the fixed point of the government's and the agents' problems as a "stackelberg game". A key difficulty arises from the fact that tracing-testing introduces heterogeneity in the population on who was tested when. Our approach allows us to conclude that time since the last test is a sufficient statistic for the agent's choice of how much to social distance. We then setup an optimal control problem for the agent and another one for the government (Sections 4 and 5 respectively) and solve the two simultaneously using the forward-backward sweep algorithm.

Two key equations summarize the government's lockdown and agents' social distancing choices. The one for the agents pins down the static and dynamic tradeoffs from social distancing. The static trade-off is described as follows: The fraction of social distancers (or the probability of social distancing) is directly proportional to the lump-sum cost of getting infected and infected others normalized by the flow cost of social distancing. Further, the dynamic trade-off adds to this proportionality the net present value of the normalized cost from getting infected while participating in social activities. Analogously, the government implements the lockdown while trading off drop in current output with mortality and net present value of future output. Both these equations also incorporate the heterogeneity of information introduced by testing.

A secondary contribution of the paper is to calibrate the model using data on testing and deaths in the United States due to Covid-19 (Section 6). It is by now well understood that the standard SIR model is poorly identified with the typical time series on the number of infected and dead (see [Fernández-Villaverde and Jones \[2020\]](#) and [Korolev \[2021\]](#)). We take a pragmatic approach by fixing the medical parameters using the aggregated wisdom of various studies

¹Section 2 details the main ingredients of the model. Section 3 describes the system of equations that quantify the propagation of the disease. Section 4 introduce behavioral response and the solves for the agent's optimization problem. And, Section 5 states and solves for the government's optimization problem.

in medical journals and then use the data to estimate three types of parameters; (i) prevalence or matching of the susceptible with the infected, (ii) efficacy of tracing and testing, and (iii) cost for agents from getting infected and infecting others. These provide a quantitative sense respectively of how rapidly the disease was spreading in the US, how effective was the government in tracing and testing the infected and asymptomatic, and how agents evaluated their decisions to social distance.

For testing we feed the model with data on daily tests that have been conducted in the US during the pandemic. By tracing we mean the efficacy to identify those in the population who are infected but haven't developed severe symptoms yet— this is captured by a parameter which is calibrated using data. Our calibrated estimate of it suggests that the US did reasonably well (on average) over the course of the pandemic in identifying and isolating the infected. Its main shortcoming was in the total number of tests available early in the pandemic say in comparison to South Korea.

How should we systemically think about the role of social distancing? We evaluate it in terms of a daily output produced by the agent, which we take approximately to be her/his daily wage. In that sense, the daily or flow cost of not partaking in any social activities is about 22% of the daily wage. In addition the agent also suffers a lump-sum disutility from getting infected; this represents the psychological cost of contracting the virus, medical costs, and the potential (probabilistic cost) of death. The calibration exercise pegs this to be approximately worth one and a half year of wages. A further 50 days of wages is the lump-sum (altruistic) disutility from potentially infecting others.

After calibrating the model, we execute the algorithm to calculate the optimal policy (Section 7). The government shuts down about 40 percent of the economic activity for a prolonged period of time, about 14 months. The lockdown is not complete in that it does not hit the upper bound (of 70 percent) at any point, but it is consistent. For the same time frame, the agents cut back around 50 percent of their social activities. So the government expects the agents to social distance which allows it to impose a less than severe lockdown and since the government locks down some part of the economic activity the agents cut back on some but not all social activities.

The underlying conceptual tack on these policy choices is that they ensure *the effective R_t , i.e. the reproduction number is maintained almost constant throughout the pandemic at one*. This controls the spread of the virus while maintaining some economic and social activities. The constancy of effective reproduction number is also in contrast to what actually transpired in the US between March 2020 and August 2021—multiple waves in infections and deaths. In fact, the total number of deaths predicted by the optimal policy is less than half of actual number for the United States as of September 2021—around 350,000 as opposed to more than 800,000 observed in the data as of writing this draft.

The optimal policy exercise described here is conducted for a specific (widely accepted) Pareto weight on total output and mortality or equivalently for a specific value of life. We then also present the Pareto frontier to illustrate the policy mix of options available with the government. For each possible value of the weight, both optimal control problems have to be solved from scratch yielding distinct value of final output and deaths. This Pareto frontier is given by the solid black line in Figure 1—net present value of total economic output is on the y -axis and total number of survivors is on the x -axis.

The dotted-blue line in Figure 1 shows the Pareto frontier when tracing-testing is shutdown, and lockdown and social distancing are the only instruments for fighting the pandemic, the curve shrinks almost uniformly. In more detailed policy experiments in the paper (in Section 8.1) we show that if the tracing-testing was more effective and aggressively pursued since the inception of the pandemic, the Pareto frontier shifts out away from the black line and the gains are quite significant. Therefore, for improved tracing-testing technology, keeping the number of survivors fixed, total output expands, and keeping total output fixed, total mortality goes down.

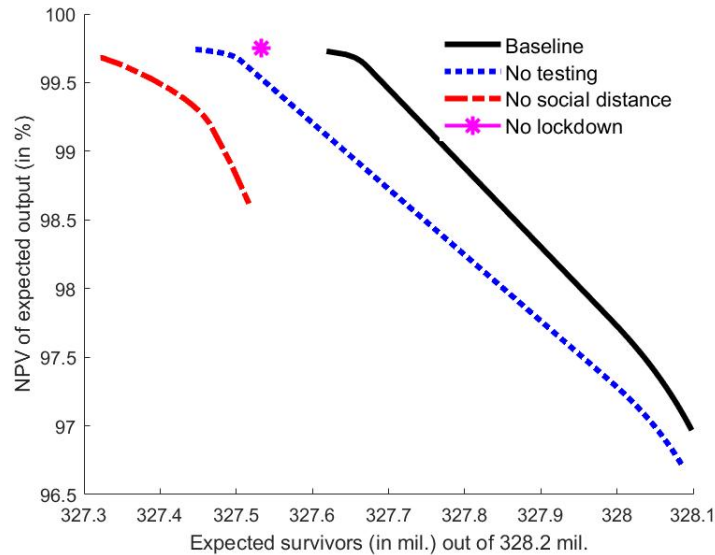


Figure 1: The Pareto frontier for different specifications of the model

What if the government imposes no lockdown? The purple star in Figure 1 represents the outcome on total output and mortality. The infection runs through the population very quickly with a single large peak during which agents social distance almost maximally. The total number of deaths is around 750k, which is significantly higher than the optimal policy and in line with the actual realized number as the pandemic has played out in the US. There are of course some economic gains since no part of the economy has been closed. We show however (Section 8.3) that even with no lockdown the government can mitigate the extent of the pandemic with more effective and aggressive testing, bringing down the cumulative number to deaths to less than 400k, while maintaining the economics gains. A provocative thought is that even with no lockdown the total number of deaths is in ballpark of the actual realized deaths seen in the data. The main difference is the spread of deaths over time. This raises the question—what is the primary goal of the lockdown policy? One interpretation of our analysis is that lockdown (as eventually practiced in reality, not in the optimal policy of our model) simply slowed down the pandemic so that the institutions fighting it could cope.²

In the next policy experiment, we ask what happens if we (hypothetically) shut down the channel of social distancing. One way to achieve this is to make the flow cost from social distancing arbitrarily high and the lump sum costs from getting infected to be low. The Pareto frontier from this exercise is given by the dashed-red curve in Figure 1. The set of feasible output and mortality attainable by the government shrinks considerably. In Section 8.3, we show that the pandemic runs through the population quickly, leading again to almost 800k deaths by the end of July 2020. The government locks down for a few months but then opens up the economy completely for the pandemic has already run amok. That is why the output doesn't fall much even though the number of deaths is large. This illustrates the significance of incorporating behavioral response in an otherwise mechanical model of disease dynamics.

In the final policy experiment, we allow the government to have greater control over restricting social interactions (Section 8.2). This changes the dynamics of infections and deaths dramatically. The idea being that autocratic governments or communitarian (as opposed to individualist) societies are able to control to social interactions of their citizenry to a much larger degree. Unsurprisingly, the deaths decrease non-linearly as we pass on greater control to a centralized authority in the social realm.

²This is echoed somewhat in Atkeson [2021b]: "this model forecast that if efforts to slow transmission were applied early but were only temporary, this dramatic first peak would be delayed but not prevented: cases and deaths would explode again once efforts to slow transmission were relaxed."

Related literature. In the rapidly growing literature on economics of epidemiology various complimentary papers look at some combination of the three forces we seek to model. Table 1 lists some of the studies done by economists and categorizes them on the basis of policy instruments.

	lockdown	tracing-testing	social distancing
Alvarez, Argente, and Lippi [2021] Acemoglu, Chernozhukov, Werning, and Whinston [2021]	✓	✓	
Bergera, Herkenhoff, Huang, and Mongey [2021] Chari, Kirpalani, and Phelan [2021]		✓	
Farboodi, Jarosch, and Shimer [2021]	✓ (exogenous)		✓
Toxvaerd [2019, 2020], Dasaratha [2020], Engle et al. [2021]			✓

Table 1: A sample of recent and related economic models of epidemiology

Broadly these papers fall within the realm of augmented SIR models, where one or more instruments from lockdown, tracing-testing and social distancing are added to lend realism to the otherwise mechanical set of equations driving disease dynamics. Lockdown and testing are typically added as a planning problem for the government and social distancing through a strategic or agency model of behavior. To the best of our knowledge, no paper so far has studied a model that incorporates all three instruments together, which is the goal of this paper. As the reader goes further, we hope to convince her/him that this is both technically challenging task and a qualitatively important one, as suggested already in Figure 1. For more detailed list of references, see McAdams [2021] for an excellent survey on the recent progress in augmented SIR models.

The importance of explicitly modeling behavior in SIR models has especially been emphasized. For example, Atkeson [2021b] writes: "behavior turns what would be a short and extremely sharp epidemic into a long, drawn out one." This also builds on a great body of work in epidemiology, where the importance of behavioral responses to improve the precision of the predictive power of the standard SIR-type model has been emphasized. For example, writing in the Proceedings of the National Academy of Sciences, Fenichel et al. [2011] state:

Results indicate that including adaptive human behavior significantly changes the predicted course of epidemics and that this inclusion has implications for parameter estimation and interpretation and for the development of social distancing policies. Acknowledging adaptive behavior requires a shift in thinking about epidemiological processes and parameters.³

To this discussion in particular, we add the nuance of how a uniform lockdown policy interacts with behavioral response, exploring their substitutability. In addition, tracing-testing introduces (i) heterogeneity in behavioral response because social distancing now is a function of the day of the latest test, and (ii) piecewise substitutability with both lockdown and social distancing in a highly non-linear fashion.

There also has been a rapidly burgeoned literature in the macroeconomics of COVID-19. Equilibrium models of the form DSGE-meets-SIR, where representative agents make labor and consumption decisions have been proposed by Eichenbaum, Rebelo, and Trabandt [2021a], Jones, Philippon, and Venkateswaran [2021], and Krueger, Uhlig, and Xie [2020], amongst others, and further extended to heterogenous agents by Kaplan, Moll, and Violante [2020] and Glover, Heathcote, Krueger, and Rios-Rull [2021], amongst others. Closer to our analysis Eichenbaum, Rebelo, and Trabandt [2021b] argue that lockdown type interventions prolong the pandemic to buy time for the health infrastructure, and there is further synergies of this force with testing and quarantining.

³See also Perra et al. [2011] and Rizzo et al. [2014] for other approaches to modeling behavior in epidemiology.

In non-SIR perspectives, [Guerrieri, Lorenzoni, Straub, and Werning \[2021\]](#) explore demand deficiencies created by the pandemic through supply side shortages. Also, [Caballero and Simsek \[2021\]](#) provide a model of asset prices spiral when the economy is hit by severe supply shock. In contrast to these, we look at a strategic framework to pin down the behavioral response of agents in an augmented SIR-type model where governments incorporate these in determining optimal policy.⁴

In terms of data work in the realm of the economics of epidemiology, our work is related to [Fernández-Villaverde and Jones \[2020\]](#) and [Korolev \[2021\]](#)—the former fits the standard SIR model to data from various parts of the world to estimate primarily the prevalence parameter and the latter argues that a unique identification of the SIR model is actually impossible. We proceed by fixing medical parameters informed by medical studies and calibrate the prevalence, tracing-testing and behavioral parameters. To the best of our knowledge, this is the first paper that tries to tease out the parameters driving tracing-testing and behavioral response by fitting the model to the data on deaths and tests.

2 Model

States and transitions. There is a continuum of identical agents with mass one who interact in discrete time indexed by $t = 1, 2, \dots$. The agents are forward-looking and discount the future with a factor $\delta_A \in [0, 1)$. At any point, a representative agent can find herself in one of five possible health states (or compartments):

- S** *Susceptible*, in this state the agent is non-infectious and non-immune;
- I** *Infected*, in this state the agent is infectious but is either asymptomatic or mildly symptomatic;
- H** *Hospitalized*, in this state the agent is infectious and clearly symptomatic, thus she is hospitalized or being treated in isolation at home;
- R** *Recovered*, in this state the agent is non-infectious and immune; it is an absorbing state;
- D** *Dead*, in this state the agent has succumbed to the infection and is dead; again an absorbing state.

The susceptible (**S**) have not had the virus, do not have immunity and may get infected in the future by coming in contact with an infected agent. The agents who get infected (**I**) are at first either asymptomatic or mildly symptomatic. This assumption is especially crucial to study the Covid-19 pandemic since one of the key difficulties has been the seemingly large number of asymptomatic carriers.⁵ Without an external intervention such as tracing/testing, agents in state **I** could potentially be indistinguishable from those in state **S**.

The infected can transition to two states: they either start showing symptoms and be hospitalized (**H**) or they can recover (**R**). The hospitalized (**H**) can also transition to two states: they can either recover (**R**) or they can die (**D**). Recovery and death are absorbing states. While death is obviously an absorbing state, assuming recovery to be absorbing implies that the agents develop immunity to the virus. Note that we are assuming that being hospitalized is a necessary step before death. This is done for simplicity. We could have added another state for non-hospitalized deaths due to Covid-19. The state **H** should be broadly interpreted as those showing clear and/or severe symptoms.

⁴There also has been recent work on understanding the socio-economic implications of the large pandemics, and take lessons that can inform policy decisions in dealing with the health and economic crisis, see [Oscar, Singh, and Taylor \[2020\]](#), [Correia, Luck, and Verner \[2020\]](#) and [Barro, Ursua, and Weng \[2020\]](#). See the survey [Brodeur, Gray, Islam, and Bhuiyan \[2021\]](#) for more details.

⁵For example, Germany's leading Covid-19 expert informing policy, Christian Drosten, had stated in the summer of 2020: "We now have evidence that almost half of infection events happen before the person passing on the infection develops symptoms – and people are infectious starting two days prior to that," ([The Gaurdian \[2020\]](#))

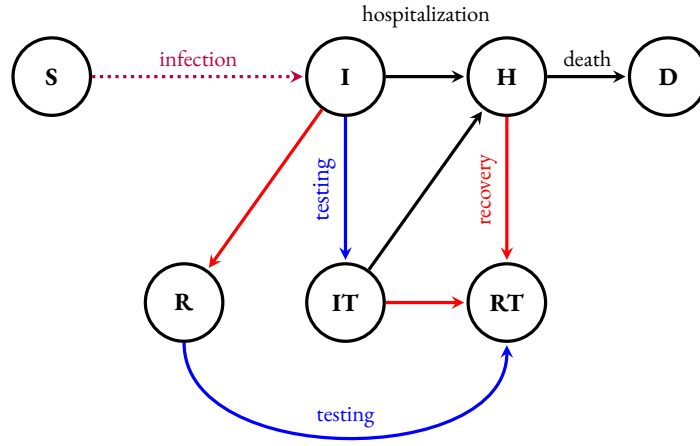


Figure 2: States and transitions in disease dynamics

Tracing and Testing. In order to separate the infected (and asymptomatic or mildly symptomatic) from the susceptible, we introduce *tracing* and *testing*. We assume that in each period a fraction of individuals who are in states **S**, **I** and **R** are traced and tested. The test perfectly reveals whether the individual is susceptible, infected or recovered. Thus, we are assuming the availability of both the viral and antibody test. This further splits **I** into two compartments, those *infected and asymptomatic* and those *infected and asymptomatic but known to be so*. These two compartments are denoted by **I** and **IT**, respectively. Similarly, **R** is split into two compartments— **R** for the *unknown recovered*, and **RT** for the *known recovered*. Figure 2 summarizes the various states and transition possibilities. Abusing notation slightly, denote by $S_t, I_t, IT_t, R_t, RT_t, H_t$ and D_t the aggregate fractions of agents in respective states at time t .

Testing is modeled through the following technology. The count of available tests is exogenous and specified by \bar{X}_t . At the end of period t a fraction τ_t is uniformly chosen amongst the set of “eligible” agents, so that exactly \bar{X}_t tests are conducted. The set of “eligible” agents includes the asymptomatic infected (**I**) and a fraction $\gamma \in [0, 1]$ of the susceptible (**S**) and unknown recovered (**R**). The parameter γ pins down the efficiency of tracing. If $\gamma = 0$, tracing is perfect for only infected people are tested, and if $\gamma = 1$, then testing is blind or completely random. Finally, tracing and testing are assumed to be history-independent, so that all agents in states $\{\mathbf{S}, \mathbf{I}, \mathbf{R}\}$ have the same likelihood of getting traced or tested, irrespective of the past state of tracing and testing.

This technology is of course simplified and buys us considerable tractability. Most natural extensions, such as making anti-body tests unavailable, making \bar{X}_t a choice variable, and introducing history dependence in testing, can be accommodated.

Two worlds. All agents engage in two types of activities in each period: *economic and social*. In each activity, the agents are randomly matched to each other. Matching is bilateral and each pair of agents can get matched independently of matching outcomes of other agents. The matching probabilities are assumed to be β_w and β_s for work (or economic) and social activities, respectively. Note that $\beta_w + \beta_s$ is equivalent to the “total prevalence” in the SIR framework, but in the context of our model it refers to the level of interaction an agent has in each period.

For social interactions, we have in mind activities such as going to a park, or visiting each others’ homes that may not directly produce any output. Moreover, the agents can control this rate of matching, as will be formalized later, by practicing social distancing. For economic activities, we have in matches that are generated while engaging in work that directly produces output. These can be controlled by government through lockdown.

Economic activity and lockdown. In addition to introducing tracing and testing, the planner also implements a *lockdown*, modeled as the stoppage of a fraction $1 - \lambda_t$ of all economic activities, bounded by $\lambda_t \geq \bar{\lambda}$.⁶ The total output of the economy during the pandemic, Y_t at time t , is assumed to have a simple structure:

$$Y_t := \lambda_t (S_t + I_t + R_t + RT_t).$$

The functional form above states that total output is equal to the total number of agents participating in economic activity, and that at any given point in time only those in states $\{\mathbf{S}, \mathbf{I}, \mathbf{R}, \mathbf{RT}\}$ are productive. We are implicitly assuming all those in states $\{\mathbf{IT}, \mathbf{H}\}$ are quarantined and do not participate in economic activities. A selective lockdown policy could allow only those in state \mathbf{RT} to work, but we think that is unrealistic, because society still needs essential services to go on, and especially given what we currently see in lockdown policies all over the world.

Social activity and behavioral response. The salient piece of the model is the introduction of behavioral response in the agents' social activities. We allow the agents to voluntarily *social distance* themselves. They have strict preferences to participate in social activities sans the virus. To capture this idea, we suppose that social distancing is costly, its flow costs are denoted by $\frac{c}{2}(1 - \alpha)^2$, where $\alpha \in [0, 1]$ is a fraction of social activities in which an agent participates. The costs are convex, more specifically quadratic, which captures an intensive margin of the agents' social distancing decision. Since the agents are forward-looking, they account for the current and future costs from social distancing. In addition to $\frac{c}{2}(1 - \alpha)^2$, the agents also face two other (lump-sum) costs: a threat of getting infected and altruistic concerns of not infecting others; they suffer disutility of ϕ^+ and ϕ^- , respectively, from the two scenarios.⁷

Vaccine. Finally, we also allow for the possibility of discovering a *vaccine* which can cure the virus. For simplicity, upon its arrival all infected agents are assumed to be cured, thus the epidemics effectively stops. Suppose the vaccine arrives at a (random) time T . The total output after the end of epidemics at $t > T$ equals to the measure of the agents who survive, that is $1 - D_T$.

We choose to model the arrival of vaccine through a negative binomial distribution. This is empirically relevant for at least two reasons. First, the distribution is parametrized by two variables: its mean $\mathbb{E}[T]$ and variance $\mathbb{V}[T]$. So, we can control both aspects of the distribution separately to reflect the reality on vaccine consensus all through the year 2020 before a credible claim on the existence of a vaccine was made. Second, the negative binomial is a sort of "repeated Bernoulli" distribution. Each period there is a probability of success or failure, and only after an exogenously specified number of success do we deem the event "vaccine developed" realized. Thus, there is a minimal time till which no vaccine can be developed, which is precisely given by $\frac{(\mathbb{E}[T])^2}{\mathbb{E}[T] + \mathbb{V}[T]}$. In what follows, we will let p_t be the probability that vaccine arrives at time t , that is $T = t$, and let $q_t := 1 - \sum_{s \leq t} p_s$ be the probability that the vaccine hasn't arrived till time t .⁸

Tying it all together. Our main interest lies in understanding how the tradeoff between economic well-being and

⁶We have in mind the necessity of essential services such as hospitals and supply chains for food and gas, and $\bar{\lambda}$ represents that threshold.

⁷The idea behind the three parameters (c, ϕ^+, ϕ^-) is that one level we want to differentiate the opportunity cost of not partaking in social activities from the cost of actually getting infected, and at another level we want to decompose the lump-sum cost of getting infected into personal costs and altruistic concerns of infecting others.

⁸Other leading candidate is the Poisson or geometric arrival, which is stationary, and hence would ensure that likelihood of the arrival of vaccine two months from the start of the pandemic is the same for it to be developed in 14 months at the 12 month mark. This distribution is a limiting case of the negative binomial in which the mean and variance coincide. This has been used, for example, by [Alvarez et al. \[2021\]](#) and [Farboodi et al. \[2021\]](#).

mortality is shaped by the simultaneous interaction of three forces: social distancing, tracing and testing, and lockdown. At the outset, the planner commits to the testing and lockdown policies which will be in effect until the development of cure. Then, within each period the following sequence of events occurs:

1. The agents in $\{\mathbf{S}, \mathbf{I}, \mathbf{R}, \mathbf{RT}\}$ participate in economic and social activities, thus creating matches and spreading the virus.
2. The agents who were in $\{\mathbf{I}, \mathbf{IT}, \mathbf{H}\}$ may recover or transition to a follow-up compartment.
3. A fraction $\tau_t \gamma$ of $\{\mathbf{S}, \mathbf{I}, \mathbf{R}\}$ is tested according to the technology described above, and amongst the tested, those in states \mathbf{I} and \mathbf{R} transition respectively to \mathbf{IT} and \mathbf{RT} .
4. The vaccine maybe discovered at time T , which is a random variable.

Finally, the (ex ante) payoff of a (representative) agent is given by $Y_A - C$, where Y_A is the total expected output, which is computed using the agent's discount factor, and C is the total cost. One way to think about Y_A is that the agent directly cares about (or is compensated for) the amount of output produced. The total cost is driven by three parameters: c , ϕ^+ and ϕ^- , flow cost from social distancing, and one time disutilities from getting infected and infecting others respectively. The payoff of the government is given by $Y_G + \xi \cdot (1 - D)$, where Y_G is the same as Y_A except we will use a (potentially) different discount factor $\delta_G \in [0, 1)$ for the government, $1 - D$ is the expected number of survivors, and ξ is the relative Pareto weight the government puts on them. Since the functional forms of these objects require some further notation, we will make them precise later.

3 Dynamics of infection

In this section we define the set of equations that governs the dynamics of transitions amongst $\{\mathbf{S}, \mathbf{I}, \mathbf{IT}, \mathbf{R}, \mathbf{RT}, \mathbf{H}, \mathbf{D}\}$. Importantly, we shut down the behavioral channel for now and try to understand mechanical aspects of the dynamics, bringing in the effect of behavioral response to the dynamics in the next section. Two types of parameters will be introduced here: the average time it takes for an agent to transfer from one state to another, and conditional on the transfer away from one state, the fraction that goes to each of the possible destination states.

Recollect that the government can control the spread of the disease by uniformly locking down a fraction of $1 - \lambda_t$ of productive agents, with the maximal lockdown capped at $1 - \bar{\lambda}$. Suppose that the susceptible and infected agents always participate in the social activities, then the stock of susceptible evolves as

$$S_{t+1} = S_t - \beta_w \lambda_t^2 S_t I_t - \beta_s S_t I_t, \quad (\Delta S_t)$$

where we count away the number of new infections due to matches at work and social activities. This is of course *one* way of modeling economic and social interactions within the standard epidemiology framework. As a first pass and for tractability we consider this additively separable assumption where the outflow of new infections is "substitutable" with coefficients β_w and β_s between the two primary activities for the agents in society.

At the inception, there is an initial seed of infection $I_0 = e_0$, and all others are susceptible: $S_0 = 1 - e_0$. Periodic matches at work and in social activities create new infections. We assume that infected agents either develop clear symptoms (and need to be hospitalized) or recover at a fixed rate on average time t_i . Taking into an account the attrition from \mathbf{I} and addition to \mathbf{IT} due to testing, we get that the following equations describe the dynamics of infected cases:

$$I_{t+1} = (1 - \tau_t) \left(\left(1 - \frac{1}{t_i}\right) I_t + \beta_w \lambda_t^2 S_t I_t + \beta_s S_t I_t \right), \quad (\Delta I_t)$$

$$IT_{t+1} = \left(1 - \frac{1}{t_i}\right) IT_t + \tau_t \left(\left(1 - \frac{1}{t_i}\right) I_t + \beta_w \lambda_t^2 S_t I_t + \beta_s S_t I_t \right). \quad (\Delta IT_t)$$

Recollect that the agents can transition from being infected, (**I**) or (**IT**), to either developing severe symptoms and being hospitalized (**H**) or recover (**R**). We assume that conditional on transitioning away from **I** or **IT**, the states **H** and **R** are reached with probabilities of transitions given by m_i and $1 - m_i$, respectively. Further, we assume that those hospitalized (**H**) transition away from that state on average time t_h . Thus, we have:

$$H_{t+1} = \left(1 - \frac{1}{t_h}\right) H_t + \frac{m_i}{t_i} (I_t + IT_t). \quad (\Delta H_t)$$

Next, the state **D** is an absorbing state, which keeps count of the number of fatalities amongst those agents that become hospitalized (**H**), and do not recover (**R**). We assume that the hospitalized agents, recover with the probability $1 - m_h$ and die with the complimentary probability m_h . Thus, we have:

$$D_{t+1} = D_t + \frac{m_h}{t_h} H_t. \quad (\Delta D_t)$$

The state **RT** is absorbing while the state **R** is “almost” absorbing in the sense that the event consists agents that have recovered without ever showing symptoms, and if tested this will be revealed and they would transfer to the state **RT**. There are three sources chipping into **R** and **RT**: infected and not tested, infected and tested, and hospitalized. These stocks evolve as follows:

$$R_{t+1} = (1 - \tau_t \gamma) \left(R_t + \frac{1 - m_i}{t_i} I_t \right), \quad (\Delta R_t)$$

$$RT_{t+1} = RT_t + \frac{1 - m_i}{t_i} IT_t + \frac{1 - m_h}{h_t} H_t + \tau_t \gamma \left(R_t + \frac{1 - m_i}{t_i} I_t \right). \quad (\Delta RT_t)$$

To sum up, the dynamics of infections and economic activity are jointly determined by the system of equations $(\Delta S_t) - (\Delta RT_t)$. In the next section we will add behavioral responses to this otherwise mechanical system. Some comments about the set of states we choose to model disease dynamics are instructive. The basic framework here is based on the celebrated SIR model which constitutes three states, namely susceptible, infected and recovered, starting at least from [Kermack and McKendrick \[1927\]](#), see also [Neher et al. \[2020\]](#) for a recent treatment. We enrich the setup by separating the infected and recovered into two categories on the basis of testing, and creating a separate state for those that show severe symptoms and thus need to be hospitalized. These two choices allow us to (i) introduce testing in a realistic way, (ii) match the widely reported stylized fact that asymptomatic carries are the largest source of contagion, at least for Covid-19, and (iii) generate at least three time series (IT_t, H_t, D_t) that are observable in the data coming out from various countries.

4 Behavioral response and the agent’s problem

We augment the epidemiological system of equations described above by giving the agents the agency to social distance themselves. This decision is motivated by two types of parameters: the flow cost of social distancing (c) and the (one

time) disutility from either getting the infection or infecting some other person (ϕ^+ and ϕ^- , respectively). All three are assumed to be fixed *behavioral* parameters. Further, the agent chooses a probability of social distancing at each time. Recollect that behavioral response is only relevant for agents in the states $\{\mathbf{S}, \mathbf{I}, \mathbf{R}\}$ because the others are either quarantined, hospitalized, known recovered, or dead.⁹

4.1 Modified dynamics

In Section 3 the stock values of all the seven states completely characterized the dynamics. However, testing introduces heterogeneity in the model which is then propagated through the agents' behavioral response. As a consequence, aggregate stock variables are no longer sufficient to keep track of the disease dynamics. In fact, we need to trace the evolution of the susceptible (\mathbf{S}), infectious (\mathbf{I}) and recovered (\mathbf{R}) agents conditional on the *time of last test*. That the time of last test is a sufficient statistic for the dynamics follows from our assumption that testing perfectly reveals an agent's current health state. After taking a test the agent is certain if she is still susceptible, infectious or recovered. In the last two scenarios the agent will transit to the follow-up compartments and never be "eligible" for testing in the future. As a result, those who are tested at time t and found to be susceptible have identical (degenerate) beliefs independently of the exact time instance t .¹⁰

Define by S_t^k , I_t^k and R_t^k the aggregate fractions of agents who are susceptible, infected and recovered, given calendar time t and time of the last test $k \leq t$. For example, an agent in state I_{30}^{35} was tested at the 35th day of the pandemic, found to be susceptible, and has since acquired the virus at some point in the last 15 days, but is asymptomatic. Here, $k = 1$ tracks the subpopulation of agents who never got any evidence, and S_t^t is the fraction of people who have just been tested to be susceptible. Further, let α_t^k be the *probability that the agent who received the last test at time k participates in social activities at time t* . We note that the total number of susceptible, asymptomatic-infected and recovered satisfy the following:

$$S_t = \sum_{k \leq t} S_t^k, \quad I_t = \sum_{k \leq t} I_t^k, \quad R_t = \sum_{k \leq t} R_t^k.$$

Similarly, define the total number of susceptible, asymptomatic-infected and recovered among those who participate in social activities as

$$\widehat{S}_t := \sum_{k \leq t} \alpha_t^k S_t^k, \quad \widehat{I}_t := \sum_{k \leq t} \alpha_t^k I_t^k, \quad \widehat{R}_t := \sum_{k \leq t} \alpha_t^k R_t^k.$$

Here, \widehat{S}_t , \widehat{I}_t and \widehat{R}_t are equilibrium objects and potentially smaller than the total numbers of people in respective compartments, viz. S_t , I_t and R_t , respectively.

When the agents make social distancing decisions, the mechanical system $(\Delta S_t) - (\Delta R T_t)$ has to be adjusted to incorporate the fact that socially inactive agents cannot spread the virus. Slightly abusing notations, we redefine certain equations taking into an account the agents' behavioral response:

⁹Note that ϕ^+ is the disutility an agent suffers from getting infected, which we take here to include the psychological cost from infecting close friends and family, and the actual perceived cost of dying or the death of these family or friends. On the other hand ϕ^- refers to the disutility the agent suffers from infecting a person he/she is randomly matched within a workspace, buying groceries or eating out or socializing in crowded places.

¹⁰The assumption that testing is history-independent, that is all agents in $\{\mathbf{S}, \mathbf{R}\}$ and all agents in \mathbf{I} have the same likelihoods of getting tested, can be relaxed to allow for conditioning on the time of last test.

$$S_{t+1}^k = (1 - \tau_t \gamma) \left(S_t^k - \beta_w \lambda_t^2 S_t^k I_t - \beta_s \alpha_t^k S_t^k \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta S_t^k)$$

$$S_{t+1}^{t+1} = \tau_t \gamma \left(S_t - \beta_w \lambda_t^2 S_t I_t - \beta_s \widehat{S}_t \widehat{I}_t \right), \quad (\Delta S_t^{t+1})$$

$$I_{t+1}^k = (1 - \tau_t) \left(\left(1 - \frac{1}{t_i} \right) I_t^k + \beta_w \lambda_t^2 S_t^k I_t + \beta_s \alpha_t^k S_t^k \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta I_t^k)$$

$$IT_{t+1} = \left(1 - \frac{1}{t_i} \right) IT_t + \tau_t \left(\left(1 - \frac{1}{t_i} \right) I_t + \beta_w \lambda_t^2 S_t I_t + \beta_s \widehat{S}_t \widehat{I}_t \right), \quad (\Delta IT_t)$$

$$R_{t+1}^k = (1 - \tau_t \gamma) \left(R_t^k + \frac{1 - m_i}{t_i} I_t^k \right) \text{ for } k = 1, \dots, t. \quad (\Delta R_t^k)$$

In addition, of course, $I_{t+1}^{t+1} = 0$, and $R_{t+1}^{t+1} = 0$. The modified dynamics makes it clear that we additively separate the two worlds of economic and social decisions, and allow lockdown policies to affect the former and social distancing to influence the latter. Testing in turn, endogenously affects both decisions by reducing those partaking in both worlds in the state **I** and potentially increasing those in the state **R**.

A quick sanity check on the system can be done by noticing that since the total population remains constant (including counting the number of deaths), the sum of all states at all points must equal unity, and in fact we have: $S_t + I_t + R_t + IT_t + RT_t + H_t + D_t = 1$.

It follows that it is without loss to ignore the equation for the state **D**, i.e., at any point in time D_t can be uniquely determined from the other state variables using the above identity. To save on notations in the future, we stack together the state variables and corresponding equations as

$$\begin{aligned} \Pi_t &:= \left((S_t^k)_{k \leq t}, (I_t^k)_{k \leq t}, IT_t, (R_t^k)_{k \leq t}, RT_t, H_t \right), \\ \Delta \Pi_t &:= \left((\Delta S_t^k)_{k \leq t+1}, (\Delta I_t^k)_{k \leq t}, \Delta IT_t, (\Delta R_t^k)_{k \leq t}, \Delta RT_t, \Delta H_t \right). \end{aligned}$$

Equations $(\Delta \Pi_t)$ jointly define the disease dynamics in an equilibrium, when the tests follows the process \overline{X}_t , the government sets the lockdown policy λ_t and the agents best-respond by choosing social distancing $(\alpha_t^k)_{k \leq t}$ as a function of the last time of being tested. In the parlance of dynamic optimization theory, these state equations are termed *forward* because they are pinned down by the initial conditions at the outset, i.e., $S_1^i = 1 - e_0$ and $I_1^i = e_0$.

4.2 Agent's optimization problem

We now turn to the agent's problem. Since a representative agent is infinitesimal, she cannot influence the aggregate state variables. Instead, the agent takes the dynamics of aggregate states as well as the government's lockdown and testing policies as given, and maximizes her own expected payoff by choosing the rate of social distancing $(\alpha_t^k)_{k \leq t}$. In this quest the agent internalizes how her current social distancing decision will affect likelihood of getting infected in the future. It turns out that the agent's problem is a dynamic control problem with the set of constraints which resembles Equations $(\Delta \Pi_t)$ modulo the fact that certain aggregate variables are taken to be fixed.

To formally define the agent's optimization problem, we should distinguish between an agent's state vector $\pi_t := ((s_t^k)_{k \leq t}, (i_t^k)_{k \leq t}, it_t, (r_t^k)_{k \leq t}, rt_t, h_t)$ and its aggregate counterpart Π_t . For example, s_t^k is the ex-ante probability that a representative agent is susceptible at time t when she last time got tested at time k . The remaining variables in π_t are interpreted similarly. We will make use the following shorthand notations: $y_t := \lambda_t (s_t + i_t + r_t + rt_t)$, and

$$s_t := \sum_{k \leq t} s_t^k, \quad i_t := \sum_{k \leq t} i_t^k, \quad r_t := \sum_{k \leq t} r_t^k, \quad \widehat{s}_t := \sum_{k \leq t} \alpha_t^k s_t^k, \quad \widehat{i}_t := \sum_{k \leq t} \alpha_t^k i_t^k, \quad \widehat{r}_t := \sum_{k \leq t} \alpha_t^k r_t^k.$$

In addition, note that the probability of being dead d_t at time t follows from the accounting identity: $s_t + i_t + r_t + it_t + rt_t + h_t + d_t = 1$.

We now have all the concepts and notation in place to define the agent's preferences. The (ex ante) expected payoff of the representative agent is given by $Y_A - C$, where

$$Y_A := (1 - \delta_A) \sum_{t=1}^{\infty} \delta_A^{t-1} \left(y_t q_t + \frac{\delta_A}{1 - \delta_A} (1 - d_{t+1}) p_t \right),$$

and

$$C := (1 - \delta_A) \frac{c}{2} \sum_{t=1}^{\infty} \delta_A^{t-1} \left((it_t + h_t + d_t) q_t + \frac{\delta_A}{1 - \delta_A} d_{t+1} p_t + \sum_{k \leq t} (1 - \alpha_t^k)^2 (s_t^k + i_t^k + r_t^k) q_t \right) + (1 - \delta_A) \sum_{t=1}^{\infty} \delta_A^{t-1} \left(\phi^+ \left(\beta_w \lambda_t^2 s_t I_t + \beta_s \widehat{s}_t \widehat{I}_t \right) + \phi^- \left(\beta_w \lambda_t^2 S_t i_t + \beta_s \widehat{S}_t \widehat{i}_t \right) \right) q_t.$$

The functional form of Y_A is straightforward, it simply aggregates the total (expected) output produced by the agent using her individual discount factor. One way to think about it is that the agent directly cares (or is compensated for) the amount of output produced.

The total cost is driven by three parameters: c , ϕ^+ and ϕ^- . The agents incur a flow cost $\frac{c}{2}(1 - \alpha_t^k)^2$ from social distancing— these arise in the states $\{\mathbf{IT}, \mathbf{H}, \mathbf{D}\}$ at their maximal level of $\frac{c}{2}$, because complete social distancing here is a necessity, and in the states $\{\mathbf{S}, \mathbf{I}, \mathbf{R}\}$, because social distancing here is a choice based on information of last test.

In addition, the agent suffers the one-shot disutility from getting infected ϕ^+ and another one from infecting others ϕ^- . The former represents the probabilistic cost of potentially getting very sick, or even facing death, and infecting members of family, while latter is an altruism parameters that captures the cost of infecting a random person at work or in the social meetings. We want the reader (relatively speaking) to think of c as being small, ϕ^+ as being large, and ϕ^- as being somewhat intermediate.

We now describe the laws of motion of the individual state vector. These differ from the laws of motion of the aggregate state vector, because each agent individually cannot influence behavior of the others on the matching markets. We have the following system:

$$s_{t+1}^k = (1 - \tau_t \gamma) \left(s_t^k - \beta_w \lambda_t^2 s_t^k I_t - \beta_s \alpha_t^k s_t^k \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta s_t^k)$$

$$s_{t+1}^{t+1} = \tau_t \gamma \left(s_t - \beta_w \lambda_t^2 s_t I_t - \beta_s \widehat{s}_t \widehat{I}_t \right), \quad (\Delta s_t^{t+1})$$

$$i_{t+1}^k = (1 - \tau_t) \left(\left(1 - \frac{1}{t_i} \right) i_t^k + \beta_w \lambda_t^2 s_t^k I_t + \beta_s \alpha_t^k s_t^k \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta i_t^k)$$

$$i_{t+1} = \left(1 - \frac{1}{t_i} \right) i_t + \tau_t \left(\left(1 - \frac{1}{t_i} \right) i_t + \beta_w \lambda_t^2 s_t I_t + \beta_s \widehat{s}_t \widehat{I}_t \right), \quad (\Delta i_t)$$

$$r_{t+1}^k = (1 - \tau_t \gamma) \left(r_t^k + \frac{1 - m_i}{t_i} i_t^k \right) \text{ for } k = 1, \dots, t, \quad (\Delta r_t^k)$$

$$r_{t+1} = r_t + \frac{1 - m_i}{t_i} i_t + \frac{1 - m_h}{h_t} h_t + \tau_t \gamma \left(r_t + \frac{1 - m_i}{t_i} i_t \right), \quad (\Delta r_t)$$

$$h_{t+1} = \left(1 - \frac{1}{t_h} \right) h_t + \frac{m_i}{t_i} (i_t + i_t). \quad (\Delta h_t)$$

In addition, $i_{t+1}^{t+1} = r_{t+1}^{t+1} = 0$ at all dates. Stack the above equations in a vector $\Delta \pi_t$, that is

$$\Delta \pi_t := \left((\Delta s_t^k)_{k \leq t+1}, (\Delta i_t^k)_{k \leq t}, \Delta i_t, (\Delta r_t^k)_{k \leq t}, \Delta r_t, \Delta h_t \right)$$

The reader can verify that Equations $(\Delta \pi_t)$ are almost identical to $(\Delta \Pi_t)$; the only difference is that we have I_t and \widehat{I}_t instead of i_t and \widehat{i}_t in the first three equations.

Finally, the agents' optimization problem is given by the choice of the social distancing $(\alpha_t^k)_{k \leq t}$ and individual state variables π_t at each date to maximize expected output net of costs subject to the family of state equations $(\Delta \pi_t)$:

$$\max_{(\alpha_t^k)_{k \leq t}, \pi_t} Y_A - C + (1 - \delta_A) \kappa P \quad \text{subject to } (\Delta \pi_t); \text{ given } (S_t, I_t, \widehat{S}_t, \widehat{I}_t) \text{ and } (\lambda_t, \tau_t).$$

In the objective, $\kappa > 0$ is a small number and P is a *punishment* term to avoid boundary solutions,¹¹ that is

$$P := \sum_{t=1}^{\infty} \delta_A^{t-1} \sum_{k \leq t} \left(\ln \alpha_t^k + \ln (1 - \alpha_t^k) \right) q_t.$$

The model admits as a special case the scenario where the agents are myopic: this will constitute substituting $\delta_A = 0$. The government, as we will model later, will always be assumed to be forward looking. This assumption of myopia for the agents dramatically simplifies all the calculations and some have argued generates a model with nonetheless reasonable predictions. As mentioned in the introduction, we do not take a call on the agent's forward looking capacity and allow the analyst the flexibility of varying the agents' discount factors as she/he deems fit to the situation being studied. To provide intuition and to point out the implications of the myopic model, we will refer sometimes to the special case as we go along.

4.3 Equilibrium response and adjoint equations

Solving for the (representative) agent's optimal social distancing rule involves setting up the Lagrangian function. Let the Lagrange multipliers corresponding to Equations $(\Delta \pi_t)$ at time t be Υ_t , i.e.,

¹¹Interior solutions for $(\alpha_t^k)_{k \leq t}$ make the execution of the algorithm that finds the optimal solution to the government's problem tractable. For context, we actually assume κ to be 10^{-12} in our numerical calculations.

$$\Upsilon_t := \left((\mathbb{S}_t^k)_{k \leq t+1}, (\mathbb{I}_t^k)_{k \leq t}, \mathbb{I}\mathbb{T}_t, (\mathbb{R}_t^k)_{k \leq t}, \mathbb{R}\mathbb{T}_t, \mathbb{H}_t \right).$$

We calculate the agent's first order-condition with respect to the control variable $(\alpha_t^k)_{k \leq t}$ and then take the first-order condition with respect to each state variable in the optimization problem to get the agent's adjoint equations.

The agent's optimal social distancing decision is thus characterized as a unique solution to the necessary first-order condition of the Lagrangian with respect to α_t^k , which is of course sufficient for optimality whenever κ is small enough:

$$\begin{aligned} 1 - \alpha_t^k + \frac{\kappa}{s_t^k + i_t^k + r_t^k} \left(\frac{1}{\alpha_t^k} - \frac{1}{1 - \alpha_t^k} \right) &= \frac{\phi^+}{c} \frac{s_t^k}{s_t^k + i_t^k + r_t^k} \beta_s \widehat{I}_t + \frac{\phi^-}{c} \frac{i_t^k}{s_t^k + i_t^k + r_t^k} \beta_s \widehat{S}_t + \\ &+ \frac{(1 - \tau_t \gamma) \mathbb{S}_t^k + \mathbb{S}_t^{t+1} \tau_t \gamma - (1 - \tau_t) \mathbb{I}_t^k - \tau_t \mathbb{I}\mathbb{T}_t}{c(1 - \delta_A) \delta_A^{t-1} q_t} \frac{s_t^k}{s_t^k + i_t^k + r_t^k} \beta_s \widehat{I}_t. \end{aligned} \quad (FOC - \alpha_t^k)$$

One intuitive way of thinking about this condition is to first consider the myopic case. Suppose $\delta_A = 0$, then the second line in Equation $(FOC - \alpha_t^k)$ disappears.¹² Now, for this myopic case, take an agent who got evidence of being susceptible at time k , then at date t she assigns probabilities

$$f_t^k := \frac{s_t^k}{s_t^k + i_t^k + r_t^k} \quad \text{and} \quad g_t^k := \frac{i_t^k}{s_t^k + i_t^k + r_t^k}$$

to the events that she is susceptible and infectious, respectively.¹³ Thus, for κ close to zero, Equation $(FOC - \alpha_t^k)$ can be re-written as follows (whenever the solution for α_t^k is interior):¹⁴

$$\underbrace{1 - \alpha_t^k}_{\text{fraction of social distancing}} \approx \underbrace{\frac{\phi^+}{c} f_t^k \beta_s \widehat{I}_t}_{\text{cost from getting infected}} + \underbrace{\frac{\phi^-}{c} g_t^k \beta_s \widehat{S}_t}_{\text{cost from infecting others}},$$

The agent experiences disutility from getting infected and infecting others which are added up. Given the aggregate behavior, each susceptible agent gets infected with a probability $\beta_s \widehat{I}_t$, and each infectious agent infects another agent with a probability $\beta_s \widehat{S}_t$. Moreover, $\frac{\phi^+}{c}$ and $\frac{\phi^-}{c}$ represent the costs of getting infected and infecting others, respectively, normalized by the cost of not being able to socialize for one period.

Having motivated the static component of Equation $(FOC - \alpha_t^k)$, allow now for $\delta_A > 0$. The agents are forward looking, and thus internalize the fact that they will transition between states in the future. Yet, they do not internalize that their social distancing decision will affect the aggregate dynamics.

Let Δ_t^k be the expected drop in the agent's continuation value due to an infection, that is $\Delta_t^k := (1 - \tau_t \gamma) \mathbb{S}_t^k + \tau_t \gamma \mathbb{S}_t^{t+1} - (1 - \tau_t) \mathbb{I}_t^k - \tau_t \mathbb{I}\mathbb{T}_t$. Re-writing Equation $(FOC - \alpha_t^k)$ for small κ and assuming that the social distancing decision is interior, we get:

$$1 - \alpha_t^k \approx \underbrace{\dots}_{\text{static, } \delta_A = 0} + \underbrace{\frac{\Delta_t^k}{c(1 - \delta_A) \delta_A^{t-1} q_t} f_t^k \beta_s \widehat{I}_t}_{\text{dynamic cost of an infection}},$$

A higher value of α_t^k increases the likelihood that the susceptible agent gets infected by $\beta_s \widehat{I}_t$. In this case the agent will transit to it_t with the probability τ_t and i_t^k with the probability $1 - \tau_t$, thus the agent's expected future value increases

¹²This is because when $\delta_A = 0$, the dual variables are zero at all dates.

¹³If the agent has been already discovered to be infected/recovered, then the agent's social distancing decision at t will be mechanical.

¹⁴In general, as $\kappa \downarrow 0$, the solution for $1 - \alpha_t^k$ converges to $\max\{0, \min\{1, \cdot\}\}$, where “ \cdot ” stays for the right-hand side of Equation $(FOC - \alpha_t^k)$. The same is true in the more general context with $\delta_A > 0$.

in a proportion to $(1 - \tau_t)\mathbb{I}_t^k + \tau_t\mathbb{I}\mathbb{T}_t$. At the same time, the likelihood that in the next period the agent will still be susceptible decreases; by the similar logic as above the agent's expected future value decreases in a proportion to $(1 - \tau_t\gamma)\mathbb{S}_t^k + \tau_t\gamma\mathbb{S}_t^{t+1}$. Finally, as in the static case, the dynamic (lump-sum) cost is deflated by the net present value of the flow cost of social distancing.

In addition to the first-order condition with respect to the control vector $(\alpha_t^k)_{k \leq t}$, we have the system of seven partial difference equations $(\Delta\mathbf{Y}_t)$ which describes the dynamics of the adjoint variables \mathbf{Y}_t . This system is similar in spirit to Equations $(\Delta\pi_t)$, which define the dynamics of individual state variables π_t . The main difference though is that that the system is *backward*, i.e., it specifies \mathbf{Y}_t as a function of \mathbf{Y}_{t+1} and other variables at time $t + 1$. The “initial” conditions are given at “infinity”:

$$\lim_{t \rightarrow \infty} \frac{\mathbb{S}_t^k}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{I}_t^k}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{R}_t^k}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{I}\mathbb{T}_t}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{R}\mathbb{T}_t}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{H}_t}{p_t \delta_A^{t-1}} = \delta_A \left(1 + \frac{c}{2}\right) \forall k.$$

As a result, the adjoint variables \mathbf{Y}_t have to be solved backwards. We present Equations $(\Delta\mathbf{Y}_t)$ for \mathbf{Y}_t and their derivations in the appendix (see Section 10.1).

Finally, equilibrium dictates that individual behavior must be consistent with aggregate behavior at every point in time, i.e., $\pi_t = \Pi_t$, $\widehat{s}_t = \widehat{S}_t$ and $\widehat{i}_t = \widehat{I}_t$. For the rest of the analysis we will utilize these equilibrium conditions to *express all the state variables in capital fonts*.

5 The government's problem

The government's objective function is given by a weighted sum of total output, which is computed using the social discount factor, and the number of survivors: $Y_G + \xi \cdot (1 - D)$, where

$$Y_G := (1 - \delta_G) \sum_{t=1}^{\infty} \delta_G^{t-1} \left(Y_t q_t + \frac{\delta_G}{1 - \delta_G} (1 - D_{t+1}) p_t \right), \quad 1 - D := \sum_{t=1}^{\infty} (1 - D_{t+1}) p_t,$$

and the number ξ is the relative Pareto weight on $1 - D$.

The government has three control variables, these are the rates of lockdown λ_t , testing τ_t and social distancing as a social planner $(\alpha_t^k)_{k \leq t}$. Moreover, the government also incorporates the agents' optimal behavior in its decision making. Thus, the Lagrangian of the government's optimal control problem takes as constraints the state equations (forward system $(\Delta\pi_t)$, now written as $(\Delta\Pi_t)$), the agent's adjoint equations (backward system, $(\Delta\mathbf{Y}_t)$) and the agent's first-order conditions with respect to the social distancing vector $(FOC - \alpha_t^k)$.

As the last constraint, the government is assumed to be exogenously restricted in its testing capacity to \overline{X}_t . We use the actual test data to feed this time series. Making the rate of testing a choice variable, τ_t , then gives us the resource constraint:

$$\overline{X}_t = \tau_t \left(\underbrace{\gamma \left(S_t - \beta_w \lambda_t^2 S_t I_t - \beta_s \widehat{S}_t \widehat{I}_t \right)}_{\text{“eligible” (S)}} + \underbrace{\gamma \left(R_t + \frac{1 - m_i}{t_i} I_t \right)}_{\text{“eligible” (R)}} + \underbrace{\left(1 - \frac{1}{t_i} \right) I_t + \beta_w \lambda_t^2 S_t I_t + \beta_s \widehat{S}_t \widehat{I}_t}_{\text{“eligible” (I)}} \right). \quad (F - \tau_t)$$

Recall that the agents are tested at the end of period t after new matches, and hence infections, are created. Then, the fraction γ of the susceptible (**S**) and recovered (**R**) added to the pool of the infected (**I**), and each agent in this set is tested at the same rate τ_t .¹⁵

¹⁵A comment on modeling choice is in order here: We could have let \overline{X}_t be a choice variable by introducing say a convex cost of tracing and

To sum up, the government seeks to maximize $Y_G + \xi \cdot (1 - D)$ by choosing the policies λ_t , $\{\alpha_t^k\}_{k \leq t}$ and τ_t which jointly control the dynamics of the state variables Π_t and agent's adjoint variables Y_t . The government is further constrained by the resource constraint on testing and the fact that the agent's social distancing decision must be a best-response. Thus, the government's problem can be stated in a consolidated way as follows:

$$\max_{\lambda_t, (\alpha_t^k)_{k \leq t}, \tau_t, \Pi_t, Y_t} Y_G + \xi \cdot (1 - D) \quad \text{subject to} \quad (\Delta \Pi_t), (\Delta Y_t), (FOC - \alpha_t^k) \text{ and } (F - \tau_t).$$

We next describe the set of optimality conditions associated with the planner's problem and the numerical algorithm to solve the model.

5.1 Optimality conditions

We use the Lagrangian approach to solve the government's problem. We now describe the system of constraints as well as the family of Lagrange multipliers which we will attach to them.

- Recollect $\Pi_t = ((S_t^k)_{k \leq t}, (I_t^k)_{k \leq t}, IT_t, (R_t^k)_{k \leq t}, RT_t, H_t)$ satisfies Equations $(\Delta \Pi_t)$. Define the government's adjoint variables corresponding to these equations to be \bar{Y}_t , i.e.,

$$\bar{Y}_t := \left((\bar{S}_t^k)_{k \leq t+1}, (\bar{I}_t^k)_{k \leq t}, \bar{IT}_t, (\bar{R}_t^k)_{k \leq t}, \bar{RT}_t, \bar{H}_t \right).$$

Similar to the agent's problem, we denote by $\bar{\Delta}_t^k$ the expected drop in the government's continuation value due to an infection, that is

$$\bar{\Delta}_t^k := (1 - \tau_t \gamma) \bar{S}_t^k + \tau_t \gamma \bar{S}_t^{t+1} - (1 - \tau_t) \bar{I}_t^k - \tau_t \bar{IT}_t.$$

- The government internalizes that each agent will best-respond. This is incorporated in the government's problem by taking the rate of social distancing $(\alpha_t^k)_{k \leq t}$ and the agent's adjoint variables Y_t to be choice variables on their own. Of course, the government is constrained by the fact that the agent's adjoint variables Y_t must respect the adjoint equations (ΔY_t) , which are formally derived in the appendix. Denote the vector of dual variables associated with this system by $\bar{\Pi}_t$:

$$\bar{\Pi}_t := \left((\bar{S}_t^k)_{k \leq t}, (\bar{I}_t^k)_{k \leq t}, \bar{IT}_t, (\bar{R}_t^k)_{k \leq t}, \bar{RT}_t, \bar{H}_t \right).$$

- Let χ_t be the vector of Lagrange multipliers associated with the testing constraint, viz. Equation $(F - \tau_t)$.
- Finally, we define $\{\eta_t^k\}_{k \leq t}$ to be the vector of Lagrange multipliers associated with the agent's first-order condition with respect to the rate of social distancing $(\alpha_t^k)_{k \leq t}$, viz. Equation $(FOC - \alpha_t^k)$. Again, this is needed to guarantee that the rate of social distancing constitutes the agent's best-reply.

The next step is to solve for the Lagrange variables $\bar{\Pi}_t$. The reader can verify that these variables satisfy system which resembles Equations $(\Delta \Pi_t)$, yet there are several noticeable differences:

testing, but the varied pathologies of why different countries ended up with their own trajectory of total tests is deeply enmeshed in political economy which requires a separate study of its own. It seems to us more reasonable, given the complexity of the model in other dimensions, to choose the time series of tests that actually materialized.

$$\bar{S}_{t+1}^k = (1 - \tau_t \gamma) \left(\bar{S}_t^k - \beta_w \lambda_t^2 \bar{S}_t^k I_t - \beta_s \left(\alpha_t^k \bar{S}_t^k - \eta_t^k S_t^k \right) \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta \bar{S}_t^k)$$

$$\bar{S}_{t+1}^{t+1} = \tau_t \gamma \left(\widehat{S}_t - \beta_w \lambda_t^2 \bar{S}_t I_t - \beta_s \left(\widehat{S}_t - \sum_{k \leq t} \eta_t^k S_t^k \right) \widehat{I}_t \right), \quad (\Delta \bar{S}_t^{t+1})$$

$$\bar{I}_{t+1}^k = (1 - \tau_t) \left(\left(1 - \frac{1}{t_i} \right) \bar{I}_t^k + \beta_w \lambda_t^2 \bar{S}_t^k I_t + \beta_s \left(\alpha_t^k \bar{S}_t^k - \eta_t^k S_t^k \right) \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta \bar{I}_t^k)$$

$$\bar{I}_{t+1} = \left(1 - \frac{1}{t_i} \right) \bar{I}_t + \tau_t \left(\left(1 - \frac{1}{t_i} \right) \bar{I}_t + \beta_w \lambda_t^2 \bar{S}_t I_t + \beta_s \left(\widehat{S}_t - \sum_{k \leq t} \eta_t^k S_t^k \right) \widehat{I}_t \right), \quad (\Delta \bar{I}_t)$$

$$\bar{R}_{t+1}^k = (1 - \tau_t \gamma) \left(\bar{R}_t^k + \frac{1 - m_i}{t_i} \bar{I}_t^k \right), \text{ for } k = 1, \dots, t, \quad (\Delta \bar{R}_t^k)$$

$$\bar{R}_{t+1} = \bar{R}_t + \frac{1 - m_i}{t_i} \bar{I}_t + \frac{1 - m_h}{h_t} \bar{H}_t + \tau_t \gamma \left(\bar{R}_t + \frac{1 - m_i}{t_i} \bar{I}_t \right), \quad (\Delta \bar{R}_t)$$

$$\bar{H}_{t+1} = \left(1 - \frac{1}{t_h} \right) \bar{H}_t + \frac{m_i}{t_i} (\bar{I}_t + \bar{I}_t), \quad (\Delta \bar{H}_t)$$

where $\bar{I}_{t+1}^{t+1} = \bar{R}_{t+1}^{t+1} = 0$, $\bar{S}_t := \sum_{k \leq t} \bar{S}_t^k$, $\bar{I}_t := \sum_{k \leq t} \bar{I}_t^k$, $\bar{R}_t := \sum_{k \leq t} \bar{R}_t^k$ and $\widehat{S}_t := \sum_{k \leq t} \alpha_t^k \bar{S}_t^k$, $\widehat{I}_t := \sum_{k \leq t} \alpha_t^k \bar{I}_t^k$.

As usual, we stack together the aforementioned equations as

$$\Delta \bar{\Pi}_t := \left((\Delta \bar{S}_t^k)_{k \leq t}, (\Delta \bar{I}_t^k)_{k \leq t}, \Delta \bar{I}_t, (\Delta \bar{R}_t^k)_{k \leq t}, \Delta \bar{R}_t, \Delta \bar{H}_t \right).$$

We note that holding the path of state variables fixed, the Lagrange variables $\bar{\Pi}_t$ follow the same law of motion as the state variables Π_t , but adjusted proportionally to $(\eta_t^k)_{k \leq t}$. In particular, Equations $(\Delta \bar{\Pi}_t)$ are forward in the sense that $\bar{\Pi}_{t+1}$ is a function of $\bar{\Pi}_t$ and other variables at time t . The initial condition for this system is $\bar{\Pi}_1 = 0$, which should be contrasted with the state equations where we have $\Pi_1 \neq 0$, i.e., $S_1^1 = 1 - e_0 > 0$ and $I_1^1 = e_0 > 0$.

The variables $\bar{\Pi}_t$ have a natural interpretation. These variables are shadow prices that convert changes in the agent's continuation payoff to changes in the government's continuation payoff. For example, consider the first period. Recall that the agent's social distancing decision at this date depends on two static costs as well as the dynamic component, which is proportional to the expected change in the agent's value due to an infection, i.e., $\Delta_1^k = (1 - \tau_1 \gamma) \mathbb{S}_1^k + \tau_1 \gamma \mathbb{S}_1^2 - (1 - \tau_1) \mathbb{I}_1^k - \tau_1 \mathbb{I}_1$. Each of these adjoint variables can be thought as a sensitivity of the agent's first period continuation payoff with respect to a certain state at date $t = 2$. For instance, \mathbb{S}_1^1 is the sensitivity related to S_2^1 , and if the agent's payoff becomes marginally more sensitive to S_2^1 , then the agent will social distance more. As a result, the planner's payoff will change in a proportion to η_1^1 , that is the multiplier on the agent's first-order condition with respect to α_1^1 . In fact, the total adjustment of the government's profit due to the marginal change in \mathbb{S}_1^1 is exactly $\bar{S}_1^1 = -\eta_1^1 (1 - \tau_1 \gamma) \beta_s S_1^1 \widehat{I}_1$.

We now look at the problem of choosing the optimal lockdown. It turns out that the optimal choice of λ_t depends on the state variables and all Lagrange multipliers through the following:

$$\lambda_t = \arg \max_{\lambda \in [\underline{\lambda}, \bar{\lambda}]} \left\{ \left((1 - \delta_G) \delta_G^{t-1} Y_t q_t + (1 - \delta_A) \delta_A^{t-1} \bar{Y}_t q_t \right) \lambda - \left(\beta_w \sum_{k \leq t} \left(\bar{\Delta}_t^k S_t^k I_t + \Delta_t^k \bar{S}_t^k I_t \right) + \beta_w \chi_t \tau_t (1 - \gamma) S_t I_t + \beta_w (1 - \delta_A) \delta_A^{t-1} \left(\phi^+ \bar{S}_t I_t + \phi^- S_t \bar{I}_t \right) q_t \right) \lambda^2 \right\}. \quad (\text{OPT-}\lambda_t)$$

To understand Equation $(\text{OPT-}\lambda_t)$, suppose for a moment that the agents are myopic, that is $\delta_A = 0$. Then, the agent's adjoint variables Y_t , and the dual variables $\bar{\Pi}_t$ attached to them, are identically equal to zero. Equation $(\text{OPT-}\lambda_t)$ reduces to the following:

$$\lambda_t = \arg \max_{\lambda \in [\bar{\lambda}, i]} \left\{ \underbrace{(1 - \delta_G) \delta_G^{t-1} Y_t q_t \lambda}_{\text{direct output}} - \left(\underbrace{\beta_w \sum_{k \leq t} \bar{\Delta}_t^k S_t^k I_t}_{\text{direct dynamic cost of an infection}} + \underbrace{\beta_w \chi_t \tau_t (1 - \gamma) S_t I_t q_t}_{\text{direct dynamic cost propagated through testing}} + \underbrace{\beta_w (1 - \delta_A) \left(\phi^+ \bar{S}_t I_t + \phi^- S_t \bar{I}_t \right)}_{\text{indirect static lump-sum cost}} \right) \lambda^2 \right\}.$$

The first term is the discounted marginal product at date t multiplied by the probability of having the pandemic going at this date. Consider the second term and recall that $\bar{\Delta}_t^k$ is the net expected change of the planner's continuation payoff due to an additional infection at state $k \leq t$. This is then multiplied by the likelihoods of transitions from the susceptible to infected, i.e., $\beta_w S_t^k I_t$, and aggregated across all testing states. As for the third term, note that an increase in economic activities creates $\beta_w S_t I_t$ new infections in total, and these then reduce the testing probability and decrease the planner's continuation payoff in a proportion to χ_t . Finally, consider the last term and recall that the agent suffers disutility ϕ^+ from getting infected and disutility ϕ^- from infecting others. These disutilities are weighted by the probabilities of respective events, pre-multiplied by their shadow prices $\bar{\Pi}_t$ and aggregated over the time of the last time. The first three terms in the above equation can be thought as direct costs, because they are unrelated to the agent's best reply.

In the general case with $\delta_A > 0$, there are two additional terms:

$$\lambda_t = \arg \max_{\lambda \in [\bar{\lambda}, i]} \underbrace{\dots}_{\text{static, } \delta_A = 0} + \underbrace{(1 - \delta_A) \delta_A^{t-1} \bar{Y}_t q_t \lambda}_{\text{indirect output}} - \underbrace{\beta_w \sum_{k \leq t} \Delta_t^k \bar{S}_t^k I_t}_{\text{indirect dynamic cost of an infection}} \lambda^2.$$

We note that the rate of economic activities at time t affects the agent's adjoint process Y_t in two ways. First, it changes each of $\mathbb{S}_{t-1}^k, \mathbb{I}_{t-1}^k, \mathbb{R}_{t-1}^k, \mathbb{RT}_{t-1}$ symmetrically and linearly, because the agent's next period output is linear in λ_t . This gives the first new term above. Second, we have the quadratic term that captures certain indirect dynamic cost. Note that $\Delta_t^k \beta_w I_t$ is the expected change in the agent's payoff multiplied by the probability of an infection at $k \leq t$. These values are pre-multiplied by \bar{S}_t^k in order to convert them to the planner's payoff, and then they are aggregated across testing states.

The government's first order condition with respect to $(\alpha_t^k)_{k \leq t}$ and τ_t can be found in the appendix (see Equations (OPT- α_t^k), (OPT- τ_t)). In short, the first set of conditions allows us to precisely pin down the set of Lagrange variables $(\eta_t^k)_{k \leq t}$ by solving a certain linear system, whereas the first-order condition with respect to τ_t determines the multiplier on the resource constraint.¹⁶

Finally, we have the system that describes the dynamics of government adjoint variables \bar{Y}_t . We relegate this system to the appendix, because it is rather complicated (see Equations ($\Delta \bar{Y}_t$)). We note that this system, as one for Y_t , is backward with the following "initial" condition:

$$\lim_{t \rightarrow \infty} \frac{\bar{S}_t^k}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{I}_t^k}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{R}_t^k}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{\mathbb{RT}}_t}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{\mathbb{RT}}_t}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{\mathbb{H}}_t}{p_t} - \delta_G^t = \xi \forall k.$$

¹⁶A comment on optimization over testing is in order here. Although, the objective is linear in τ_t , we must have $\tau_t \in (0, 1)$ at all dates, because the process \bar{X}_t , which we infer from the actual testing data, is strictly positive (and not too large). To guarantee that $\tau_t \in (0, 1)$ is indeed optimal we set the term in the objective in front of it to zero.

5.2 Numerical algorithm

We now briefly outline the algorithm which we use to solve for the optimal policies. First of all, we truncate the problem by fixing a certain large terminal time \bar{t} such that $p_{\bar{t}}$, the probability that the vaccine has not arrived by time \bar{t} , is sufficiently small. The solution to the truncated problem is characterized by exactly the same equations as above, but the terminal conditions for two backward systems have to hold exactly at $t = \bar{t}$.

Our numerical approach is a variation of the so-called *forward-backward sweep algorithm*. The forward-backward sweep is a standard iterative method to solve an optimal control problem with control dynamics in the form of partial difference/differential equations. The idea is to start with a policy, use it to simulate state variables forward and then solve for adjoint variables backward. These two sets of variables jointly produce a new policy. If it is sufficiently close to the original one, then the algorithm stops. Otherwise, it updates the initial policy as a convex combination of the old and new ones with the weights of $1 - \varepsilon > 0$ and $\varepsilon > 0$, respectively.

We adapt the forward-backward sweep to two nested control problems: one for the agent and the other is for the government. Thus, we will have two loops. In the inner loop we solve the agent's problem with a fixed lockdown policy, whereas in the outer loop we update the lockdown policy and Lagrange multipliers on the resource constraint and agent's first-order condition with respect to the rate of social distancing.

Let λ_t for $t = 1, 2, \dots, \bar{t}$ be given. Then, the inner loop is as follows:

Step 0 Select the rate of social distancing $(\alpha_t^k)_{k \leq t}$.

Step 1 Solve for the state variables Π_t and rate of testing τ_t using Equations $(\Delta\Pi_t)$ and $(F - \tau_t)$.

Step 2 Solve for the agent's adjoint variables Υ_t using Equations $(\Delta\Upsilon_t)$.

Step 3 Find the rate of social distancing $(\tilde{\alpha}_t^k)_{k \leq t}$ using Equation $(FOC - \alpha_t^k)$.

Step 4 If the distance between $(\alpha_t^k)_{k \leq t}$ and $(\tilde{\alpha}_t^k)_{k \leq t}$ is small enough, then stop the inner loop. Otherwise, update $(\alpha_t^k)_{k \leq t}$ to $(1 - \varepsilon) \cdot (\alpha_t^k)_{k \leq t} + \varepsilon \cdot (\tilde{\alpha}_t^k)_{k \leq t}$ and go to Step 1.

The outer loop is as follows:

Step 0 Select the rate of lockdown λ_t , the Lagrange multipliers on the agent's first-order condition with respect to the rate of social distancing $(\eta_t^k)_{k \leq t}$ and the Lagrange multipliers on the resource constraint χ_t .

Step 1 Use the inner loop to obtain the rate of social distancing $(\alpha_t^k)_{k \leq t}$.

Step 2 Solve for the state variables Π_t , shadow cost variables $\bar{\Pi}_t$ and rate of testing τ_t using Equations $(\Delta\Pi_t)$, $(\Delta\bar{\Pi}_t)$ and $(F - \tau_t)$.

Step 3 Solve for both sets of adjoint variables Υ_t and $\bar{\Upsilon}_t$ using Equations $(\Delta\Upsilon_t)$ and $(\Delta\bar{\Upsilon}_t)$.

Step 4 Find the rate of lockdown $\tilde{\lambda}_t$ using Equation $(OPT-\lambda_t)$, compute the Lagrange multipliers $(\tilde{\eta}_t^k)_{k \leq t}$ and $\tilde{\chi}_t$ using Equations $(OPT-\alpha_t^k)$ and $(OPT-\tau_t)$, respectively.

Step 5 If the distance between $(\lambda_t, \{\eta_t^k\}_{k \leq t}, \tau_t)$ and $(\tilde{\lambda}_t, \{\tilde{\eta}_t^k\}_{k \leq t}, \tilde{\tau}_t)$ is small enough, then stop the inner loop. Otherwise, update $(\lambda_t, (\eta_t^k)_{k \leq t}, \tau_t)$ to $(1 - \varepsilon) \cdot (\lambda_t, (\eta_t^k)_{k \leq t}, \tau_t) + \varepsilon \cdot (\tilde{\lambda}_t, (\tilde{\eta}_t^k)_{k \leq t}, \tilde{\tau}_t)$ and go to Step 1.

The inner loop takes a few seconds to terminate, whereas the outer loop needs around ten minutes to produce the solution. Several comments are in order. First, it is well-known that the forward-backward sweep algorithm might run into cycles, especially, when ε is large. We control for this by fine-tuning this parameter, in fact, in our simulations with $\varepsilon = 0.01$ the algorithm converges monotonically. Second, we validate global optimality by multi-starting the algorithm with different initial values chosen at random. We are not aware of alternative procedures which can be used to solve for the global optima in the case of two nested control problems.

6 Calibration

Predicting disease dynamics, calculating optimal policy and evaluating policy experiments all require reliable parametrization of the model. In principle we could let the entire set of parameters be identified using time-series of data on state variables. This unfortunately runs into problems for even an ideal dataset will leave any standard SIR model unidentified (see [Korolev \[2021\]](#) and [Fernández-Villaverde and Jones \[2020\]](#)).

We take the following pragmatic approach. We fix government policies, medical and other miscellaneous parameters using the aggregated wisdom of research from medicine and social science, and calibrate behavioral parameters to match the time series of data available on the state variables. The underlying conceptual thought here is to borrow liberally from other scientific studies, but try our best to tease out the (i) *initial prevalence*—the building block of SIR framework, (ii) *social distancing costs*, and (iii) *efficacy of tracing-testing* by fitting the data on the progress of the pandemic thus far to the predictions of our model.

Let the vector of parameters that will be estimated be given by $\theta := (\beta_w, \beta_s, c, \phi_+, \phi_-, \gamma)$. In what follows we briefly describe our dataset, the set of fixed parameters and the estimation process to pin down θ . All additional details can be found in the appendix.

6.1 Method

We combine data from two sources for our estimation exercise. Time series data for daily deaths and positive cases are obtained from "[The New York Times Repository](#)", while time series data for daily hospitalizations and daily tests are sourced from "[COVID Tracking Project](#)". Our analysis covers data between February 29, 2020 ($t = 1$), when the first death was reported in the US, till February 14, 2021 ($t = 352$).¹⁷

The primary objective of the estimation exercise is to parametrize our model so that it rationalizes the US Covid-19 data obtained from the above mentioned sources. This requires us to solve the agent's optimization problem given the *actual* lockdown and testing policies followed by the government. Since there is no reliable source to estimate the actual lockdown function implemented by the US government, we use an approximation which is roughly consistent with the government response index published by [Hale et al. \[2021\]](#) (see [Oxford COVID-19 Government Response Tracker](#) for the data). Figure 3a compares our approximation of lockdown function and the Oxford GRT index.¹⁸

As for the testing capacity, we compute the test count from the state level data as opposed to simply using the publicly available aggregate data, for example, see "[COVID Tracking Project](#)". The reason is that data is collected at the state level in the US, and the states have practiced distinct ways of reporting negative tests. Some states report

¹⁷The collected data suffers from the "weekend" effects: new observations are added in chunks with lags of several days contributing to the high volatility. To mitigate this issue, we smoothen the data by constructing moving averages in the window of seven days for each variable. This smoothening of Covid-19 data is standard and has also been employed by other studies, see for example, [Fernández-Villaverde and Jones \[2020\]](#).

¹⁸The Oxford GRT reports a variety of indices to capture the policy responses undertaken by different countries. These include policies for economic reform, stringency measures to reduce transmission, healthcare reform and so on. To construct the index illustrated in Figure 3a, we use two stringency measures used by the US government to curb economic activities, namely workplace closure and stay-at-home requirements.

people tested at least once (unique people), others report either testing encounters or samples collected. To the best of our knowledge, neither "COVID Tracking Project" nor other data providers, i.e., "The New York Times Repository", explicitly have corrected for these differences in reporting.

Ideally, we would like to use testing encounters for our calibration exercise but since only a few states (mainly less populated) reported this metric imputing it might result in imprecise estimates. Another reason for confusion is that more than one test is often conducted on a patient in a short span, especially if she/he tests positive. So, we construct \bar{X}_t by imputing and aggregating the state level testing data on the number of unique people who got tested, which can be thought as a lower bound on the total US testing capacity, in-sample till February 14, 2021. Then, we extrapolate the count of available tests out-of-sample by fitting a 2nd degree polynomial function of time. Figure 3b plots the actual number of conducted tests in US and its out-of-sample extrapolation.




Figure 3: Assumed government policies

Having set the lockdown and testing policies, we now present the set of medical and other auxiliary parameters that are taken to be fixed to calibrate the other parameters of interest. Table 2 summarizes these fixed parameters. In short, we rely here on the accumulated knowledge thus far, and we discuss each parameter in details in the appendix.

Parameter	Value	Definition	Source/Target
t_i	10	Average infectious period (incl. latency)	CDC Planning Scenario
t_h	7	Average hospitalization period	
m_i	0.0176	Hospitalized as % of infected	$m_i = IFR/m_h$, where $IFR = 0.3\%$ (Ioannidis [2021])
m_h	0.1705	Daily deaths as % of hospitalized	Estimated by OLS
\bar{ET}	540	Expected day the pandemic ends	New York Times [2020]
\bar{VT}	180	Variance in the day pandemic ends	
$\bar{\lambda}$	0.3	Share of essential services	Used by multiple studies incl. Alvarez et al. [2021], Acemoglu et al. [2021] and Farboodi et al. [2021]
δ_G	0.9999	Government's discount factor	
δ_A	0.9999	Agent's discount factor	
e_o	0.007%	Share of infected at the start	23,000 infected agents given the population of 328.2 mil.

Table 2: Fixed parameters

Given the paucity and limited reliability of the Covid-19 data mentioned above, only the flow data on deaths and

positive cases are used to match the model parameters. We employ the following loss function for this purpose:

$$\omega \cdot \sum_{t=1}^{352} \left(\Delta D_t(\theta) - \Delta \tilde{D}_t \right)^2 + (1 - \omega) \cdot \sum_{t=1}^{352} \left(X_t^+(\theta) - \tilde{X}_t^+ \right)^2$$

where $\theta = (\beta_w, \beta_s, c, \phi_+, \phi_-, \gamma)$ is the vector of parameters being estimated, $\Delta D_t(\theta)$ is the predicted flow of deaths at time t , $X_t^+(\theta)$ is the predicted flow of positive tests, i.e., the sum of newly hospitalized and known infected, at date t , $\Delta \tilde{D}_t$ and \tilde{X}_t^+ are their empirical counterparts. The weights ω and $1 - \omega$ ensure that both variables get the same importance in the computation of the loss function.^{19, 20}

To avoid over-fitting the model we impose two additional restrictions. First, we require $\beta_w = \beta_s$, which is broadly in agreement with the American Time Use Survey that employed Americans (on average) spend about half time at work and half time in social activities (see Eichenbaum et al. [2021a]). Second, we fix $\phi^- = \frac{\phi^+}{10}$, i.e., we are assuming the one-time cost of infecting another person outside of immediately friends and family to be 10 percent of the cost that we would suffer from getting infected ourselves and potentially risking infection for our loved ones as well. This is broadly subjective and can be adjusted in light of more evidence.

6.2 Estimation

Recollect that we are trying to fit the following two time series: the number of daily deaths and the number of daily positive cases. The inner loop of the numerical algorithm described in Section 5.2 is able to generate the fixed point. The robustness of the fixed point is verified by multi-starting the algorithm. Figures 4a and 4b compares the model fit and data for daily deaths and daily positive cases, and Table 3 reports the calibrated parameters.

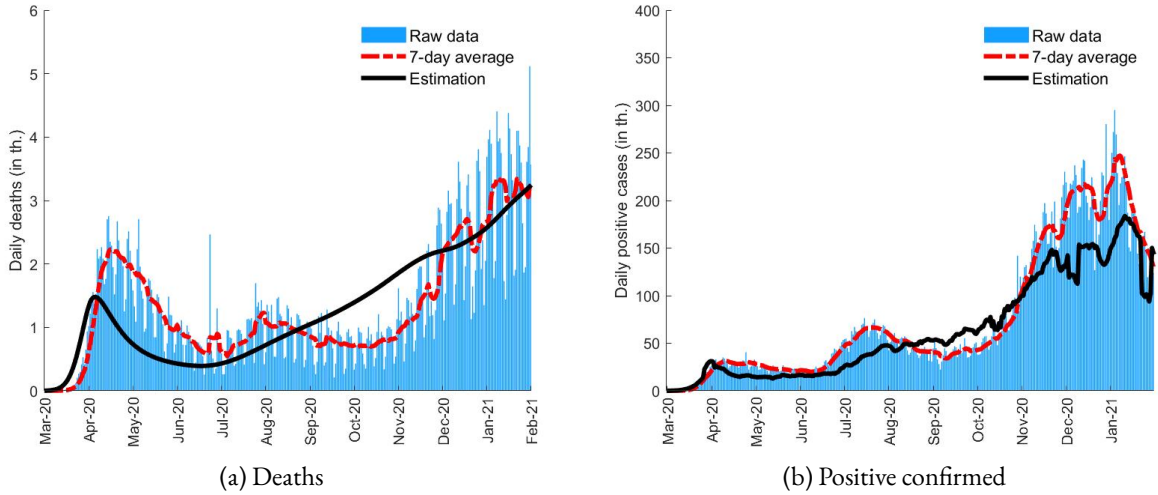


Figure 4: Model fit vs data

Overall, the model with just 4 free parameters fits the data reasonably well. It turns out that our estimates match almost perfectly the cumulative numbers of deaths and positive cases by February, 14. However, they do not pick up magnitudes of some waves smoothing them over time. This can be partially attributed to factors not present in the model such as overload of the healthcare system, information percolation about the pandemic, seasonality in contact

¹⁹This is required since the daily number of positive cases vastly outnumbers the number of daily deaths in the US. For instance, the mean daily deaths in the US was 1284 compared to the mean daily positive cases of 70982.

²⁰It is worthwhile to note that we fix the test capacity \tilde{X}_t^+ at the actual empirical level. Hence fitting $X_t^+(\theta)$ to the data ensures that the predicted flow negative cases fit the data well too.

Parameter	Value	Definition
β_w	0.2046	Prevalence in economic activities
β_s	0.2046	Prevalence in social activities
γ	0.0976	Tracing-testing efficacy
c	0.4403	Daily opportunity cost of infection
ϕ_+	523.932	Disutility of getting infected
ϕ_-	52.393	Disutility of infected someone else

Table 3: Calibrated parameters.

rates, i.e., academic year and holidays, etc.²¹

Our estimate of the total prevalence $\beta_w + \beta_s$ turns out to be 0.4092 that implies that the initial rate of growth of infection is slightly more than 0.3. In other words, the number of infected agents doubles in 3 days at the beginning of our analysis. This is in line with previous studies. For example, Farboodi et al. [2021] calibrated the total prevalence parameter to 0.43, and Eichenbaum et al. [2021a] estimated the transmission rate in social activities to be around 0.16, which is close to our estimate of $0.5 \times 0.4092 = 0.2046$.

The efficacy of trace-testing γ is estimated to be 0.0976. Recall that a fraction $1 - \gamma$ of the susceptible (**S**) and unknown recovered (**R**) is excluded from testing based on tracing, so here $\gamma = 0$ is perfect tracing technology and $\gamma = 1$ denotes completely random testing. The estimate of efficacy of trace-testing is novel to the literature, and it essentially tells us that conditional on the number of tests, the tracing and testing infrastructure in the United States has been reasonably effective in targeting the infected population. However, as we will argue later, the delay in scaling up testing had significant consequences.

For the behavioral parameters, we estimate three types of costs which each agent incurs. The daily opportunity cost of missing work and social activities due to being quarantined, measured by c , is estimated to be 0.4403. We find ϕ^+ to be estimated to be around 523.9. This represents the psychological cost from knowing that one has been infected plus the cost of facing a potential death from the virus or risking the health of a close family member. Then, ϕ^- is set to 53.4, which is the (ex ante) cost of getting someone else infected and contributing indirectly to all their direct costs. As mentioned in the introduction, one way to interpret these numbers is in terms of daily wage as the unit of measurement.

7 The optimal policy

In this section we report the results from executing the numerical algorithm described in Section 5.2 to calculate the optimal lockdown and testing policies for parameters specified in the previous section and the value of life ξ set to 20. This specific value of ξ is in line with Hall, Jones, and Klenow [2020] and has been chosen in other papers that have followed. We use this number for the simulation exercise and reporting of some of the results, but then we also map out the Pareto frontier to exposit the full set of policy options available to the government.

Figure 5a plots the optimal lockdown and average fraction of social distancers at each date over the course of 20 months. The first thing to notice is that government internalizes that individuals will voluntarily social distance at moderate levels in order to avoid getting infected. Towards the beginning of the pandemic in March 2020, agents are willing to give up around 55% of their social activities (dashed red line). Interestingly this is the highest level of social distancing that agents engage in throughout the pandemic. Over time the extent of social distancing declines very

²¹Note that we could substantially “improve the fit” visibly by freeing up some more parameters from Table 2. However, this would surely result in an overfitting of the model.

gradually till August 2021 to a 50%. As the possibility of the pandemic ending increases, the optimal social distancing levels decline rather steeply to 0 between the period of August 2021 and October 2021.²²

Given this moderate but persistent level of social distancing, the government imposes an early lockdown of around 40% which then remains relatively stable (solid black line). It rationally internalizes that the pandemic is going to spread, and under the expectation of vaccine arriving in Summer/Fall 2021 it chooses an intermediate but persistent lockdown to avoid a large outbreak of infections. The rate of lockdown declines gradually to around 35% till July 2021, followed by a rapid decline as the pandemic is expected to end. However, the government never imposes anywhere close to the maximal lockdown levels in the economy throughout the period of the pandemic.




Figure 5: Optimal lockdown and social distancing

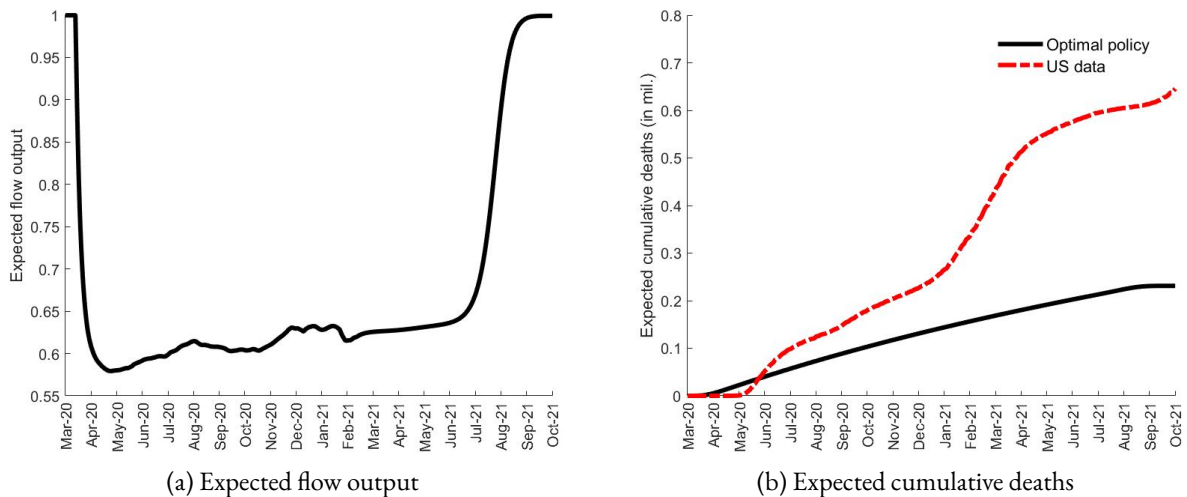


Figure 6: Flow output and cumulative deaths

The rationale behind the optimal lockdown can be seen by examining the famed reproductive number R_t . It is an approximation of the expected number of secondary cases produced by a single infected agent at a specific date t .²³

²²The small spike in social distancing before the sharp decline is due to the uncertainty in vaccine arrival or modeling end of the pandemic. Since the government starts to ease the lockdown, if the vaccine is very likely to arrive but hasn't quite arrived yet, agents may social distance more to control the infection rate.

²³A variant of this expression without testing is reported in studies like Farboodi et al. [2021] and Hall et al. [2020].

$$R_t := \underbrace{\frac{1}{1 - (1 - \tau_t)(1 - 1/t_t)}}_{\text{"avg. duration" of state I}} \cdot \underbrace{(1 - \tau_t) \left(\beta_w \lambda_t^2 S_t + \beta_s \widehat{S}_t \frac{\widehat{I}_t}{I_t} \right)}_{\text{\# of new infections}}.$$

Recall that an infected agent can transmit the disease through two channels: economic activities, in which case she/he infects $\beta_w \lambda_t^2 S_t$, and social activities, in which case she/he infects $\beta_s \widehat{S}_t \frac{\widehat{I}_t}{I_t}$ agents. Taking into an account that τ_t of these new secondary cases will be test positive immediately at the end of date t , the total number of secondary cases produced is given by $(1 - \tau_t) \left(\beta_w \lambda_t^2 S_t + \beta_s \widehat{S}_t \frac{\widehat{I}_t}{I_t} \right)$. Then, we note that this infected agent stays in **I** if and only if she/he has not moved to **IT** (tested positive), **R** (recovered) or **H** (hospitalized). Thus, if testing rate is close to being a constant, she/he is expected to spend approximately $\frac{1}{1 - (1 - \tau_t)(1 - 1/t_t)}$ days in the compartment of unknown infected agents. We call R_t the *effective reproductive number*.

Figure 5b plots the effective reproductive number implied by the optimal policy and agents' social distancing decision. The government chooses the lockdown function internalizing the agent's best response in a way that R_t is approximately 1 for almost 17 months. This in turn means that modulo the variation in total available tests the number of active cases is targeted to be approximately constant over this time frame, which can be seen by examining Equation (ΔI_t). Importance of reducing the effective reproductive number to one has been extensively addressed in public discourse, and we believe that our results provide arguments in favor of keeping this as the target of policy.

The two outcomes that define the objective function of the government are depicted in Figure 6. Recollect that the total number of survivors is weighted by the $\xi = 20$ for these calculations. Figure 6a shows the flow of output as the pandemic evolves. Total output falls almost by 45% percent initially according to the optimal policy and then gradually claws its way back to its full potential as the government eases lockdown. This graph of course reflects the lockdown policy in Figure 5a because total output in our model is simply equal to the labor supply.

Figure 6b plots the cumulative number of deaths predicted by the optimal policy (solid black line) and the actual realized number (dashed red line). The two series depart fairly quickly around the two-three month mark, and while the slope of the actual number gets a bump as winter 2020 approaches, the optimal policy ensures the slope of the number of deaths remains more or less constant rising up to about three hundred thousand deaths, less than half of the actual number of deaths seen in the data, which has already passed seven hundred and fifty thousand as of writing this draft of the paper.

Finally, going beyond the specific value of ξ , in Figure 7 the entire Pareto frontier is presented. This graph is produced by varying ξ on a grid of thirty points ($\xi = 0$ to 30), and then solving the government's problem for each possible value. Thus, the curve traces all efficient combinations of total output and mortality by changing the value of the Pareto weight ξ . That the curve is downward sloping is obvious. Interestingly, it first has an approximately constant slope followed by a concave transition. Mathematically speaking, the optimization problem here is mapping out the concave closure of the *actual* Pareto frontier. So, the initial straight part of the frontier is representative of a weakly convex frontier—hence increasing returns in terms of lives saved for every unit drop in output. Eventually the concave part diminishes these returns quite substantially.

For low values of ξ , especially $1 \leq \xi \leq 10$, the trade-off is stark: the first one percent permanent drop in output increases the expected number of survivors by 281 thousand. For larger values of ξ , for example $13 \leq \xi \leq 30$, a one percent drop in output increases the expected number of survivors by 158 thousand. Thus, even a moderate emphasis on mortality in the objective function produces significant gains in lives saved with a concomitant linear (or even lesser) decline in output. However, as the Pareto weight is increased the trade-off becomes less rewarding—the number of lives saved per unit drop in output decreases marked by the concave rate of transition.




Figure 7: Mapping the Pareto frontier on output and lives saved

8 Policy experiments

In this section we explore policy experiments by changing two parameters— efficiency of tracing and testing and the fraction of prevalence the government can directly control—and also removing lockdown and social distancing piecemeal from the optimization problem.

8.1 Improving tracing and testing

There has been substantial variation in the way different countries have approached large scale testing to fight the pandemic. Some seem to have delayed in making testing widely available and others in making contact tracing particularly effective. Here we do two experiments on the baseline.

First, we allow for early testing. As can be seen in Figure 3b, testing starts late in the pandemic, almost at the two and half month mark. We expand the total set of test capacity at each date to be its highest attained value in the data, which is one million tests a day. We call this the high testing counterfactual. Second, we allow the tracing and testing to be much more efficient. We do that by reducing γ by half—from 0.09 to 0.045. This ensures a larger proportion of infected are now being tested and quarantined. Figure 8 plots the key variables for both counterfactuals and also for the combined scenario of high and efficient testing.

The pairwise substitutability of better testing with lockdown and social distancing is immediately clear from Figures 8a and 8b respectively. When testing is scaled earlier and is more efficient the optimal extent of lockdown and social distancing diminish significantly. On average, lockdown drops by almost 40 percent, and social distancing by about 20 percent. On the social distancing dimension, if an agent is not selected for testing, she/he is relatively more confident of not being infected and thus is willing to engage more in social activities. For lockdown too, since the government is more certain that the pool of agents who are not quarantined or recovered are not infected, it needs to lockdown less.

Improved tracing-testing leads to less total infections; however, lower lockdown and social distancing lead to more infections. So the aggregate effect of improved tracing-testing on mortality is ambiguous. In fact in Figures 8c and 8d, we can see that for fixed level of testing, introducing greater efficiency keeps the number of deaths to almost the same level, in fact slightly increases it. Recall these numbers are being determined endogenously through the solutions to the

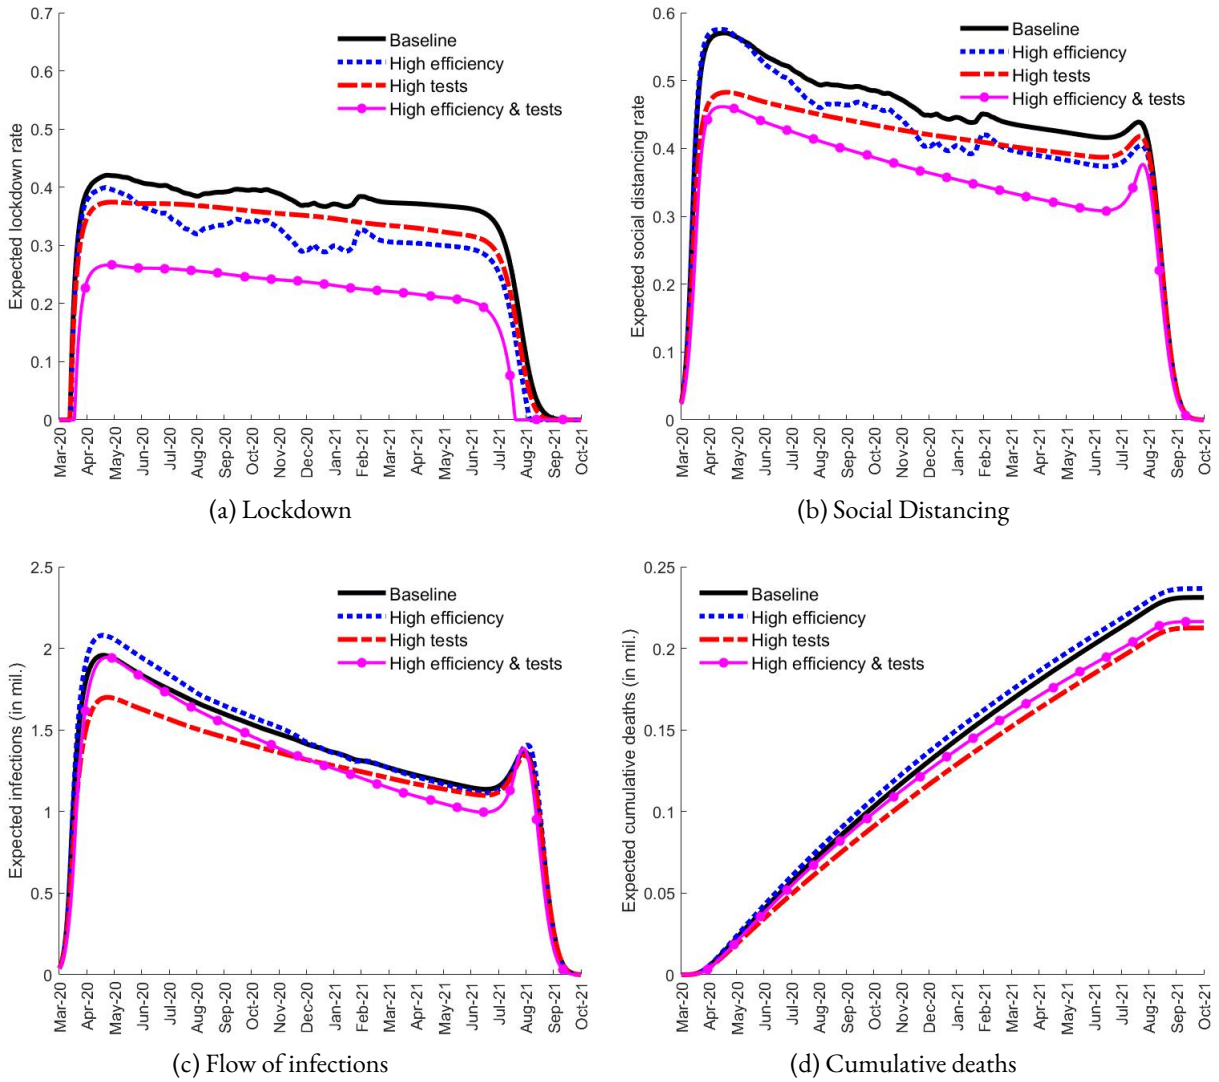


Figure 8: Counterfactuals for early testing, and effective tracing and testing

government and the agents' problems for a fixed Pareto weight. It turns out that the government is trading off some mortality for a much larger increase in output.

Figure 9 plots the Pareto frontier for four different specifications of tracing-testing. Introducing either high number of tests or higher efficiency of testing pushes out the Pareto frontier. And, introducing both considerably expands it. A simple way to think about this expansion is that for a fixed level of mortality, the government now achieves a much higher output and similarly for a fixed level of output it can now achieve considerably more survivors (or less deaths). These findings are analogous to [Acemoglu et al. \[2021\]](#)—they feature testing levels tagged to age, we have a uniform testing policy; however, agents in our model choose their level of social interactions, in theirs they are mechanical.

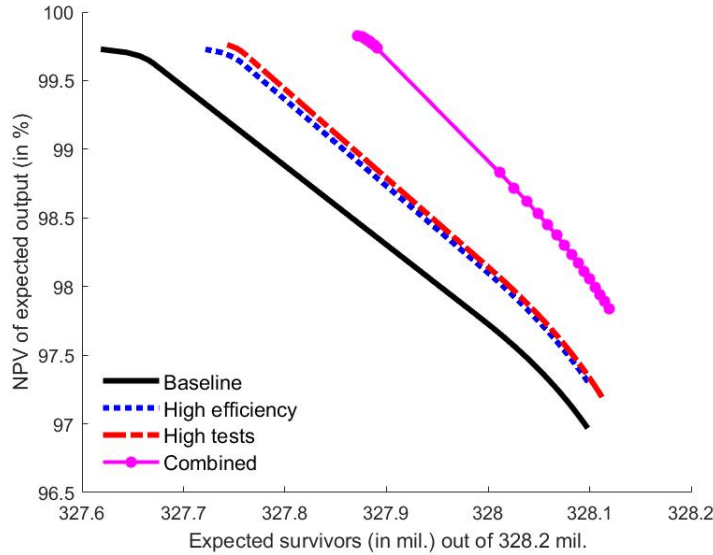


Figure 9: The Pareto frontier for high and improved testing

8.2 Greater control of the government

There has been intense discussion through the pandemic that autocracies have been able to control the diseases better than democracies (see [Narita and Sudo \[2021\]](#)). While the moral and institutional intricacies of that questions are outside the scope of this paper, we can shed some light on it through a simple policy experiment. How do the basic results change by increasing the control of the government over social interactions?

We do this through the following reduced form experiment. Recollect that $\beta_w = \eta\beta$ and $\beta_s = (1 - \eta)\beta$, where β can be considered as the aggregate prevalence parameter which sets the contagiousness of this disease, and η is the fraction of infections that arise at work and $1 - \eta$ the fraction that arise in social interactions, respectively. We started out by setting $\eta = 0.5$, which is motivated by the American Time Use Survey (see [Eichenbaum, Rebelo, and Trabandt \[2021a\]](#)). We now present the results for $\eta = 0.4, 0.6$ and 0.9 . The idea is to vary the matching intensity in the social interactions part of Equation (ΔS_t) and its subsequent avatars. So, we interpret $\eta = 0.4$ to mean that the fraction of matches within the direct control of the government decreases. And, the way we interpret $\eta = 0.6, 0.9$ is that government has been able to enforce masking and other other such policies (which decreases matches) more effectively. The agents must follow these policies to a varying degree, as η increases.

As η increases the government essentially tries to kill the pandemic through a more extensive lockdown for a larger fraction of infections matches are under its control. Why does the government do this? Remember that saving lives has a direct value in the government's objective function and every life saved produces more output in the future. Even though the flow of output decreases, its net present value is higher because of lives saved through the stricter lockdown. And since the government has greater control on matches the stricter lockdown has a greater bang for buck. The agents on the other hand social distance less because given the strict lockdown in place, the probability of them contracting the infection in social interactions is lower—they try to take advantage of whatever social interactions are permitted. Overall, greater governmental control reduces infections and total deaths. See [Figure 10](#) for the results.

8.3 No lockdown and social distancing

Several governments have faced the dilemma of whether to lockdown economic activities and by how much. Sweden is perhaps the most widely discussed case, it went in for little or almost no lockdown over the early part of the pandemic

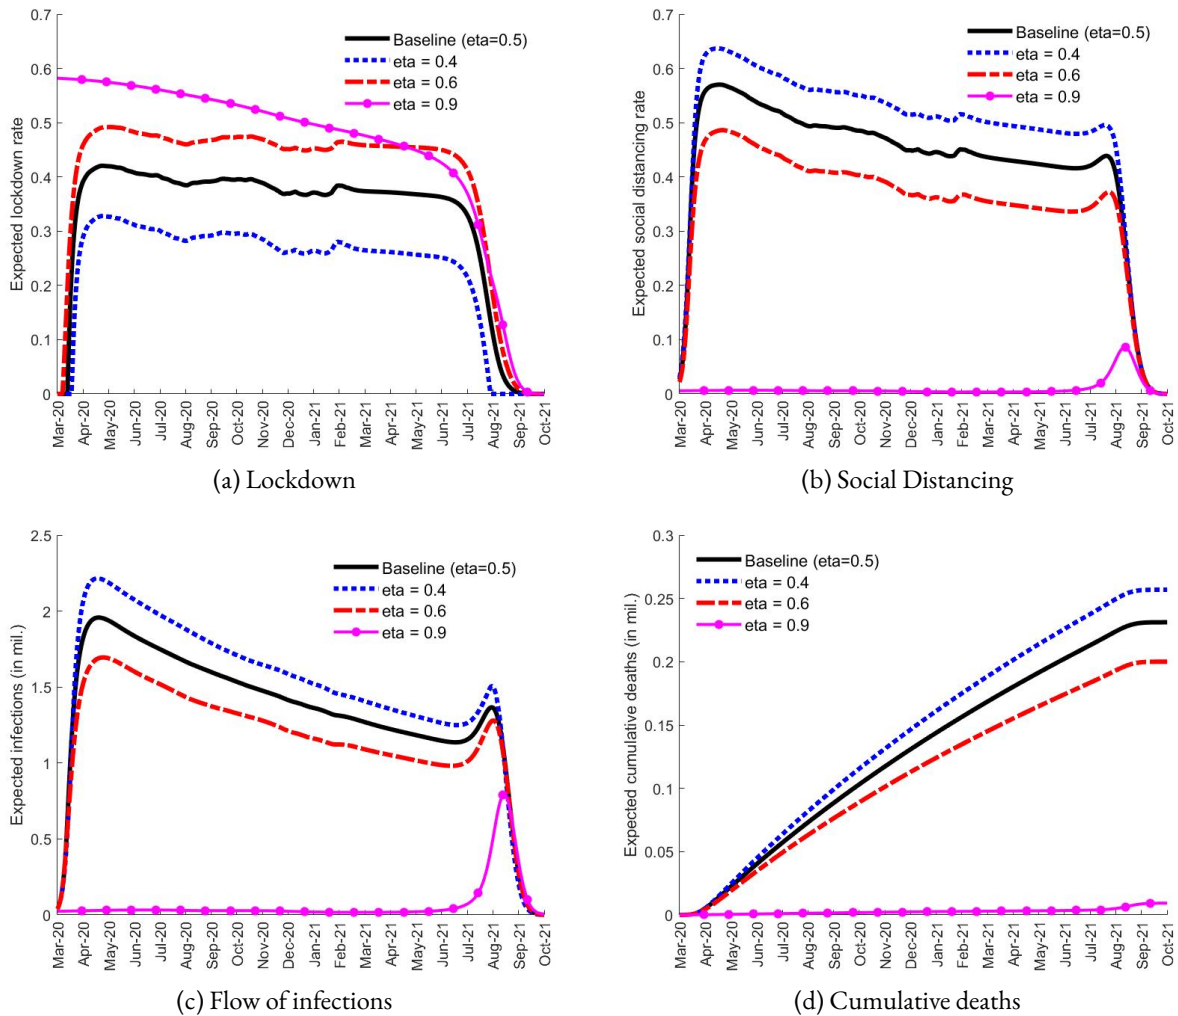


Figure 10: Counterfactuals for different levels of η , where $\beta_w = \eta\beta$ and $\beta_s = (1 - \eta)\beta$

(Financial Times [2021a]). In this section we first evaluate the policy experiment where lockdown is exogenously fixed at zero. We analyze this scenario for both the benchmark constellation of parameters and for the case where testing-tracing is more extensive and efficient.

Figures 11a and 11b show that pandemic speeds through the population with one large peak of infections at the inception in response to which the agents social distance to a peak level of shutting down 90% of all social activities (red lines in both graphs). Correspondingly, the number of deaths rises rapidly to about 750 thousands by July 2020, and then stays constant as the population reaches herd immunity (see red line in Figure 11c). This makes one ponder because the *actual* number of deaths in the US as of writing this draft is also in the range of 750 thousands. A provocative thought then is that with its current capacity in testing, the US could have ended up with the same numbers of total deaths even without a lockdown—the difference would be in the distribution of fatalities over time.²⁴

The solid purple line in Figure 11 plots social distancing, infection and deaths for the no lockdown case when it is accompanied with more extensive and effective tracing-testing. It is clear that if testing-tracing capacity is ramped up then even with no lockdown the peak of infections is much smaller, and the total number of deaths is reduced significantly, the number being lower than what we see actually realized in the data.

²⁴There are obvious caveats here: Agents need to have understood the implications of the pandemic at the outset and were forward looking in their social distancing decisions. Moreover, hospitals all over the country could have handled the extreme peak in infections shown in Figure 11b. Subject to these caveats, this point about the equivalence of deaths is made to highlight that one of the goals of optimal policy is in-fact to spread out the pandemic to allow the government and society to manage it better.

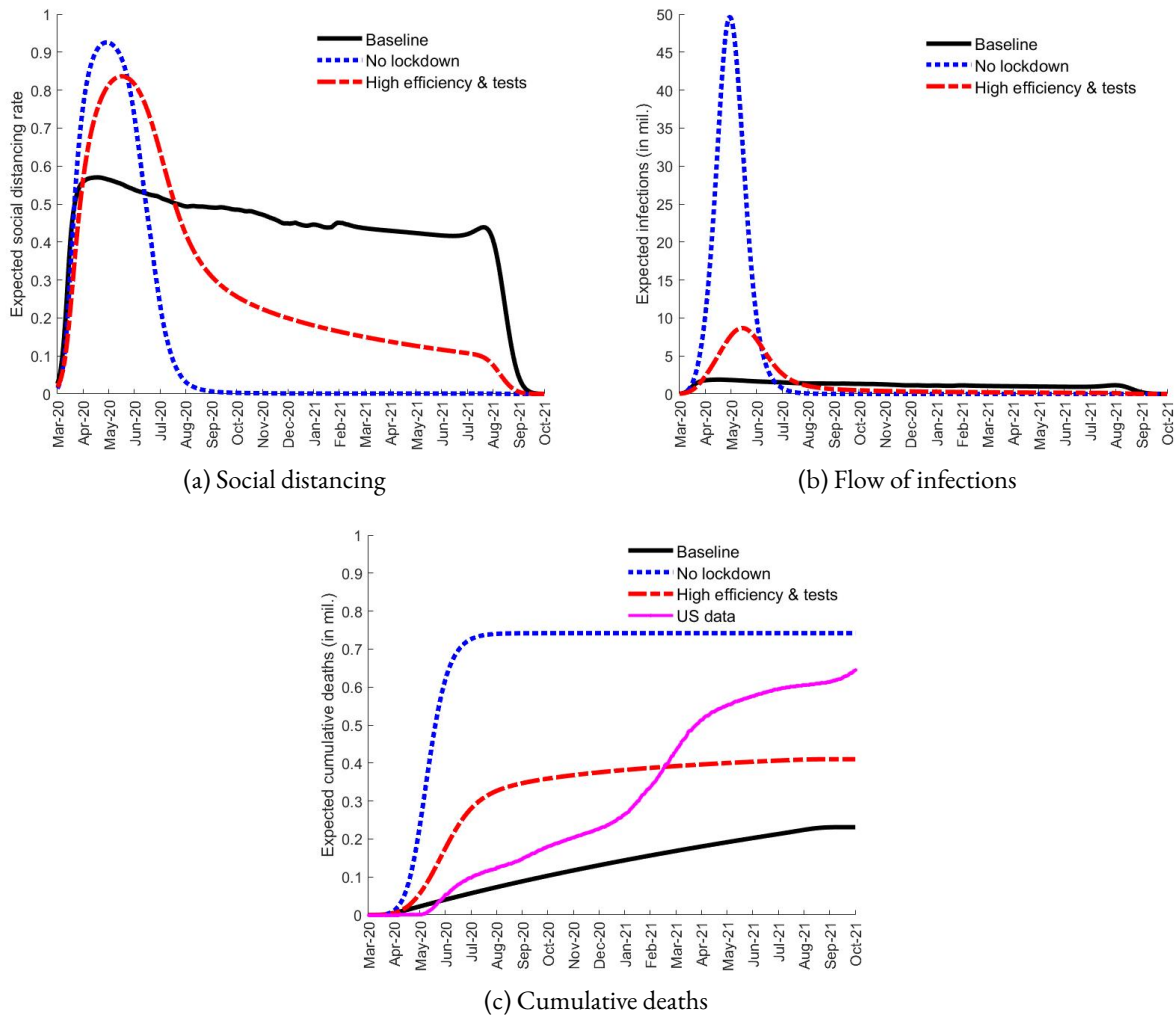


Figure 11: Counterfactuals for no lockdown

In the next policy experiment we ask the question: what if social distancing is shutdown in the model? This scenario is not meant to be taken literally, it is simply a theoretical exercise to understand what the analyst or the modeler may miss if agency or behavioral response is not directly modeled in an SIR framework. Figure 12 plots the results for this exercise. A pattern similar to the no lockdown case emerges, which highlights the complementarity of the two forces. However, things are more severe when social distancing is shut down in the sense that government locks down maximally while the infections peak (dotted blue in Figures 12a and 12b respectively). The total number of deaths rises quickly to more than 750 thousands (dotted blue line in Figure 12c). Again increasing efficacy of testing-tracing reduces the peak of infections and total number of deaths (dashed red lines in Figures 12b and 12c). The reduction in the peak is much more extreme in the no lockdown case in comparison to the no social distancing case. This is because better testing technology is used to draw out the pandemic longer through social distancing; whereas is in the no social distancing case, because of the concern for total output, the pandemic is made to run through faster. This contrast can be seen by comparing the shape of the cumulative deaths curve (dashed red line) in Figures 11b and 12c).

9 Final remarks

In this paper we set out to incorporate two policy instruments, namely lockdown and tracing-testing, into a standard epidemiology framework augmented by behavioral response. The theoretical analysis tackles the challenge of captur-

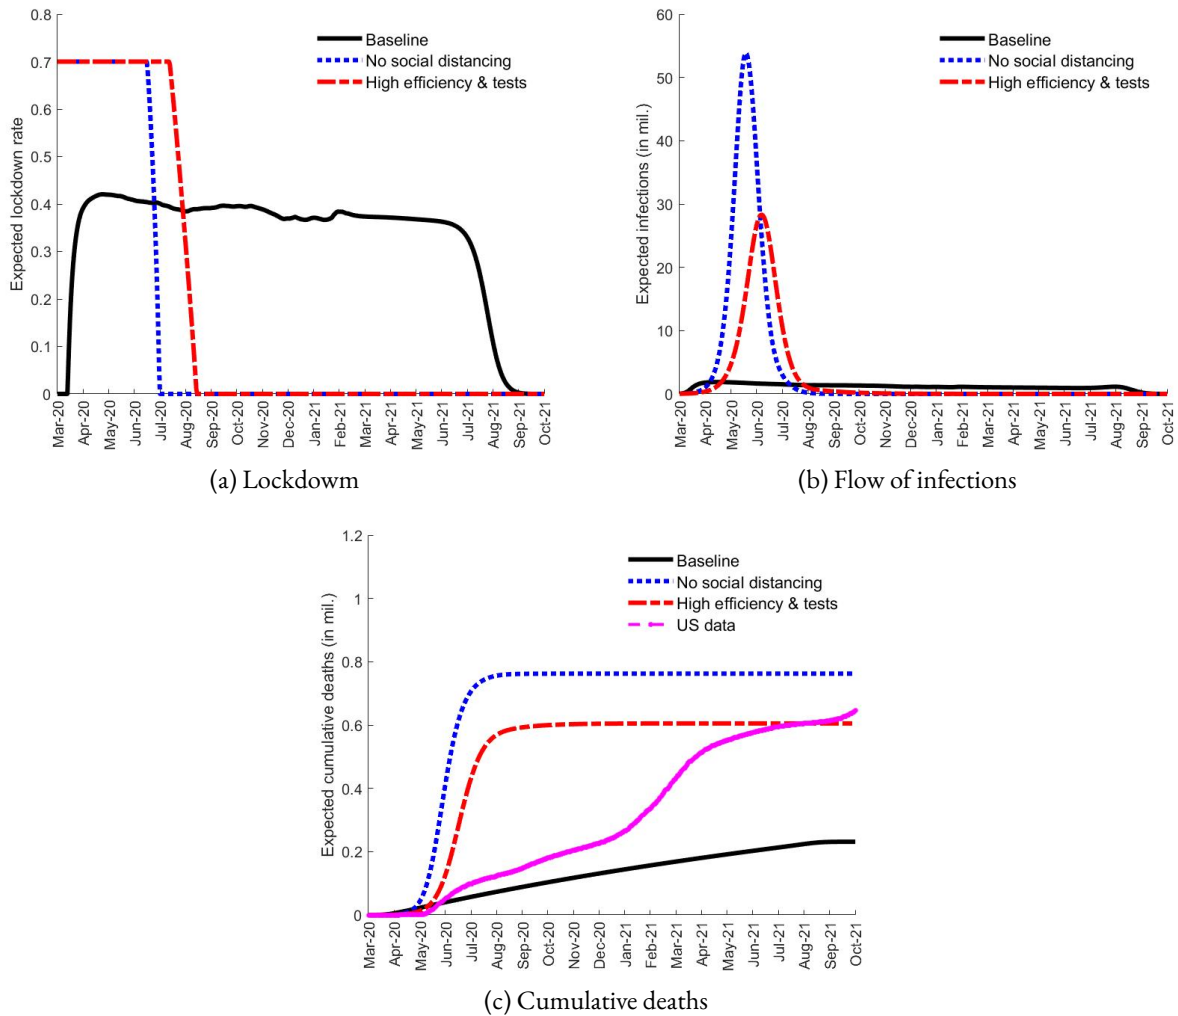


Figure 12: Counterfactuals for no social distancing

ing the pairwise substitutability of each of the three forces and the heterogeneity introduced by testing. The model is calibrated to the data on the time series of deaths and tests conducted in the US (from Feb 2020 to Feb 2021) to calculate the three key parameters—prevalence of the virus, efficiency of tracing-testing, and costs from social distancing. The testing capacity is fixed by the available data on tests conducted in the US during the same time frame. Under the optimal policy, the government locks down at a moderate but persistent rate and agents social distance too at a moderate but persistent rate, each picking up the slack of the other. These two together contribute in keeping the effective reproductive number at one throughout the pandemic. Pareto frontier maps the feasible set of economic output and total fatalities. Three different policy experiments are conducted to evaluate how the predictions of the model would change under some commonly debated policy prescriptions.

In this final section, we now briefly discuss some limitations of our analysis, and extensions of the model that could be useful in addressing these concerns.

Modeling behavioral response. We model behavioral response by introducing costs from social distancing that constitute a flow component as well as a lumps-sum component from getting infected and infecting others. This seems like a natural way to incorporate the incentives for social distancing, as has been done, for example by [Toxvaerd \[2020\]](#) and [Farboodi et al. \[2021\]](#). However, we depart from existing work by assuming that those who have recovered without showing severe symptoms cannot know so unless they are tested. There are potentially interesting micro-foundations

to these decisions of social distancing worth thinking about carefully. What role do authorities and peer effects play? This requires a better modeling of information percolation.²⁵ Also, are individualistic societies less or more likely to internalize social distancing into their daily realities than those that are community driven?

Modeling testing. We introduced a very simple tracing and testing technology wherein a fraction of susceptible, and infectious and recovered without symptoms can be tested every period. The efficacy of testing is controlled by a single parameter. Testing technology can be enriched by first separating rapid tests that only discover infection from antibody tests that also discover whether the person was infected in the past. Restricting testing to only the former, which is more realistic, would change the effectiveness of testing, but by how much is an open question, at least in the context of our model. Another avenue of future research is to focus tracing, testing and quarantining more towards the elderly, who are significantly more vulnerable.²⁶

Modeling heterogeneity in the population. It is now clear Covid-19 is much more lethal to the elderly and people with certain specific co-morbidities. One way to model this richness is to create separate compartments for different sections of the population, define system of equations for the evolution of disease dynamics in each compartment with matching functions within and across them. Work along these lines has been done by [Baqae et al. \[2020\]](#) and [Acemoglu et al. \[2021\]](#) amongst others. Marrying features of behavioral response to these frameworks or introducing age dependent compartments in our model remains an open question.

Modeling economic activity and lockdown. We modeled lockdown as a time dependent policy that shuts down a fraction of economic activity. Total output in turn is modeled simply as the total labor supply. These are of course abstractions that do capture the gist of the problem, but miss out some important ingredients. First, lockdown can be selective across various regions of a country, and in varying intensities; see [Fajgelbaum et al. \[2020\]](#) and [Giannone, Paixao, and Pang \[2020\]](#) on how this matters for relaxing lockdown policies over time. Second, drop in output itself can kill lives, especially in developing countries where daily wage workers are abundant and social safety nets are weak, see for example [Barnett-Howell and Mobarak \[2020\]](#) for a discussion. This would, amongst other things, change the way we choose a Pareto-weight for optimal policy. Third, a richer model of production could be interesting to examine. How capital and supply chains are impacted by the pandemic and governmental response to it are salient too.

Modeling the end of the pandemic. The end of the pandemic is modeled as a negative binomial distribution. The parameters are chosen so that the pandemic could not have ended 360 days after its inception and it almost surely ends 720 days after with an average duration of 540 days. In reality it is difficult to define what the "end of the pandemic" means. We gave it the interpretation of the arrival of a vaccine. To that end, a more careful modeling of the roll out of the vaccine and related factors such as vaccine hesitancy, and mutations to the virus also seems important.

On the Pareto frontier and value of life. Putting a Pareto-weight on the tradeoff between economic output (or consumption) and rates of mortality is a deep philosophical question perennially relevant to how we structure society

²⁵See [Bandyopadhyay, Chatterjee, Das, and Roy \[2021\]](#) for a model of optimal government policy where agents learn about the pandemic from the intensity of the lockdown.

²⁶In response to the question, should all countries be testing uniformly, Germany's leading Covid expert informing policy, Christian Drosten, had stated early in the pandemic: "I'm not sure. Even in Germany, with our huge testing capacity, and most of it directed to people reporting symptoms, we have not had a positivity rate above 8%. So I think targeted testing might be best, for people who are really vulnerable... This is not fully in place even in Germany, though we're moving towards it. The other target should be patients in the first week of symptoms, especially elderly patients who tend to come to hospital too late at the moment... And we need some kind of sentinel surveillance system, to sample the population regularly and follow the development of the reproduction number" ([The Gaurdian \[2020\]](#)).

and institutions, see [Schelling \[1968\]](#) for a classic reference. More recently [Hall, Jones, and Klenow \[2020\]](#) have taken up the question in the context of the recent Covid-19 crises. For our analysis, we were reluctant to pick a particular number and thus eventually provided an entire frontier. A potentially interesting question here is to measure how much societies value life through a revealed preference argument with the pandemic as a natural experiment.

On pandemic fatigues, second waves and new variants. A salient feature of the pandemic across countries has been the rise of infection in waves. While our model is flexible enough to match the data in the US which went through at least two waves of the pandemic, it does not explicitly model the underlying behavioral reasons for it. [Atkeson \[2021a\]](#) takes first steps in this direction by introducing “a one-time reduction in the semi-elasticity of the transmission rate to the daily death rate late in the year,” as pandemic fatigue. Much more work needs to be done in understanding behavioral underpinnings of the pandemic fatigue and how optimal policy should take it into account. Moreover, the successive waves of the pandemic have been associated with the rise of new “variants”. At the time of the writing of this draft the Delta variant of Covid has created much havoc and the Omicron variant is on the rise. Incorporating the ebb and flow of variants precisely because of the global nature of the spread of the virus is another challenge for modeling.

On mental health consequences. We have abstracted here from the mental health consequences of lockdown and to some extent of social distancing as well. There is widespread evidence of a spurt in mental health troubles during the course of Covid-19, see for example [COVID-19 Mental Disorders Collaborators \[2021\]](#) and [Mayo Clinic \[2021\]](#). To the best of our knowledge, these haven’t been incorporated in the economic models of the pandemic yet. It presents an important challenge for policymakers for blanket lockdowns policies may have other perverse consequences not immediately observable, which need to be incorporated in trade-offs of policymaking.

In conclusion, the most important caveat is that while we hope this paper helps economists and maybe perhaps even epidemiologists and public health researchers think carefully about the ways in which to incorporate behavioral response in models of disease dynamics and evaluate optimal policy, great care should be taken to interpret any of these insights to inform actual policy. As social scientists, we are in the long-haul of understanding socio-politico-economic contexts of this deep health crisis.

10 Appendix

10.1 Derivation of the agent’s optimal behavior

The agent’s optimization problem is described in details in [Section 4.2](#): Taking the aggregate variables $(S_t, I_t, \widehat{S}_t, \widehat{I}_t)$ and (λ_t, τ_t) as given, the agent solves

$$\max_{(\alpha_t^k)_{k \leq t}, \pi_t} Y_A - C + (1 - \delta_A)C\kappa P \quad \text{subject to} \quad (\Delta\pi_t).$$

We now set up the Lagrangian for the agent’s problem denoting the undiscounted dual variables associated with Equations $(\Delta\pi_t)$ by $\Upsilon_t = (\mathbb{S}_t^k)_{k \leq t+1}, (\mathbb{I}_t^k)_{k \leq t}, \mathbb{I}\mathbb{T}_t, (\mathbb{R}_t^k)_{k \leq t}, \mathbb{R}\mathbb{T}_t, \mathbb{H}_t)$. The Lagrangian can be unpacked as fol-

lows:

$$\begin{aligned}
& \underbrace{(1 - \delta_A) \sum_{t=1}^{\infty} \delta_A^{t-1} \left(y_t q_t + \frac{\delta_A}{1 - \delta_A} (1 - d_{t+1}) p_t \right)}_{Y_A} - \underbrace{(1 - \delta_A) \frac{c}{2} \sum_{t=1}^{\infty} \delta_A^{t-1} \left((i_t + h_t + d_t) q_t + \frac{\delta_A}{1 - \delta_A} d_{t+1} p_t + \sum_{k \leq t} (1 - \alpha_t^k)^2 (s_t^k + i_t^k + r_t^k) q_t \right)}_{\text{III}} \\
& - \underbrace{(1 - \delta_A) \sum_{t=1}^{\infty} \delta_A^{t-1} \left(\phi^+ (\beta_w \lambda_t^2 s_t I_t + \beta_s \widehat{s}_t \widehat{I}_t) + \phi^- (\beta_w \lambda_t^2 s_t i_t + \beta_s \widehat{S}_t \widehat{i}_t) \right) q_t}_{C} + \underbrace{(1 - \delta_A) c \kappa \sum_{t=1}^{\infty} \delta_A^{t-1} \sum_{k \leq t} (\ln \alpha_t^k + \ln (1 - \alpha_t^k)) q_t}_{P} + \\
& + \sum_{t=1}^{\infty} \sum_{k \leq t} \mathbb{S}_t^k \left(\underbrace{-s_{t+1}^k + (1 - \tau_t \gamma) \left(s_t^k - \beta_w \lambda_t^2 s_t^k I_t - \beta_s \alpha_t^k s_t^k \widehat{I}_t \right)}_{\Delta s_t^k} \right) + \sum_{t=1}^{\infty} \mathbb{S}_t^{t+1} \left(\underbrace{-s_{t+1}^{t+1} + \tau_t \gamma \left(s_t - \beta_w \lambda_t^2 s_t I_t - \beta_s \widehat{s}_t \widehat{I}_t \right)}_{\Delta s_t^{t+1}} \right) + \\
& + \sum_{t=1}^{\infty} \sum_{k \leq t} \mathbb{I}_t^k \left(\underbrace{-i_{t+1}^k + (1 - \tau_t) \left(\left(1 - \frac{1}{t_i} \right) i_t^k + \beta_w \lambda_t^2 s_t^k I_t + \beta_s \alpha_t^k s_t^k \widehat{I}_t \right)}_{\Delta i_t^k} \right) + \sum_{t=1}^{\infty} \mathbb{I}_t \left(\underbrace{-i_{t+1} + \left(\left(1 - \frac{1}{t_i} \right) i_t + \tau_t \left(\left(1 - \frac{1}{t_i} \right) i_t + \beta_w \lambda_t^2 s_t I_t + \beta_s \widehat{s}_t \widehat{I}_t \right) \right)}_{\Delta i_t} \right) + \\
& + \sum_{t=1}^{\infty} \sum_{k \leq t} \mathbb{R}_t^k \left(\underbrace{-r_{t+1}^k + \left((1 - \tau_t \gamma) \left(r_t^k + \frac{1 - m_i}{t_i} i_t^k \right) \right)}_{\Delta r_t^k} \right) + \sum_{t=1}^{\infty} \mathbb{R}_t \left(\underbrace{-r_{t+1} + \left(r_t + \frac{1 - m_i}{t_i} i_t + \frac{1 - m_h}{h_t} h_t + \tau_t \gamma \left(r_t + \frac{1 - m_i}{t_i} i_t \right) \right)}_{=\Delta r_t} \right) + \\
& + \sum_{t=1}^{\infty} \mathbb{H}_t \left(\underbrace{-h_{t+1} + \left(\left(1 - \frac{1}{t_h} \right) h_t + \frac{m_i}{t_i} (i_t + i_t) \right)}_{\Delta h_t} \right),
\end{aligned}$$

where we used our standard shorthand notations: $s_t = \sum_{k \leq t} s_t^k$, $i_t = \sum_{k \leq t} i_t^k$, $r_t = \sum_{k \leq t} r_t^k$, $\widehat{s}_t = \sum_{k \leq t} \alpha_t^k s_t^k$, $\widehat{i}_t = \sum_{k \leq t} \alpha_t^k i_t^k$, $\widehat{r}_t = \sum_{k \leq t} \alpha_t^k r_t^k$, $y_t = \lambda_t (s_t + i_t + r_t + r_t)$ and $d_t = 1 - (s_t + i_t + i_t + r_t + r_t + h_t)$.

The optimal level of social distancing is characterized by the necessary first-order condition ($FOC - \alpha_t^k$). For completeness, we copy the condition:

$$1 - \alpha_t^k + \frac{\kappa}{s_t^k + i_t^k + r_t^k} \left(\frac{1}{\alpha_t^k} - \frac{1}{1 - \alpha_t^k} \right) = \frac{\phi^+}{c} \frac{s_t^k}{s_t^k + i_t^k + r_t^k} \beta_s \widehat{I}_t + \frac{\phi^-}{c} \frac{i_t^k}{s_t^k + i_t^k + r_t^k} \beta_s \widehat{S}_t + \frac{\Delta_t^k}{c(1 - \delta_A) \delta_A^{t-1} q_t} \frac{s_t^k}{s_t^k + i_t^k + r_t^k} \beta_s \widehat{I}_t,$$

where $\Delta_t^k = (1 - \tau_t \gamma) \mathbb{S}_t^k + \tau_t \gamma \mathbb{S}_t^{t+1} - (1 - \tau_t) \mathbb{I}_t^k - \tau_t \mathbb{I}_t$. It is easy to see that the equation above admits a unique solution for small enough $\kappa > 0$.²⁷

The remaining first-order conditions with respect to the vector of state variables are as follows:

$$\begin{aligned}
\mathbb{S}_t^k &= (1 - \tau_{t+1} \gamma) \mathbb{S}_{t+1}^k + \tau_{t+1} \gamma \mathbb{S}_{t+1}^{t+2} - \Delta_{t+1}^k \left(\beta_w \lambda_{t+1}^2 I_{t+1} + \beta_s \widehat{I}_{t+1} \alpha_{t+1}^k \right) + \delta_A^t \left(1 + \frac{c}{2} \right) p_t + \\
& + (1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \left(1 - (1 - \alpha_{t+1}^k)^2 \right) - \phi^+ \left(\beta_w \lambda_{t+1}^2 I_{t+1} + \beta_s \alpha_{t+1}^k \widehat{I}_{t+1} \right) \right) q_{t+1}, \quad k = 1, \dots, t+1 \quad (\Delta \mathbb{S}_t^k) \\
\mathbb{I}_t^k &= \left(1 - \frac{1}{t_i} \right) \left((1 - \tau_{t+1}) \mathbb{I}_{t+1}^k + \tau_{t+1} \mathbb{I}_{t+1} \right) + \frac{1 - m_i}{t_i} \left((1 - \tau_{t+1} \gamma) \mathbb{R}_{t+1}^k + \tau_{t+1} \gamma \mathbb{R}_{t+1} \right) + \frac{m_i}{t_i} \mathbb{H}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t + \\
& + (1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \left(1 - (1 - \alpha_{t+1}^k)^2 \right) - \phi^- \left(\beta_w \lambda_{t+1}^2 S_{t+1} + \beta_s \alpha_{t+1}^k \widehat{S}_{t+1} \right) \right) q_{t+1}, \quad k = 1, \dots, t \quad (\Delta \mathbb{I}_t^k) \\
\mathbb{I}_t &= \left(1 - \frac{1}{t_i} \right) \mathbb{I}_{t+1} + \frac{1 - m_i}{t_i} \mathbb{R}_{t+1} + \frac{m_i}{t_i} \mathbb{H}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t, \quad (\Delta \mathbb{I}_t) \\
\mathbb{R}_t^k &= (1 - \tau_{t+1} \gamma) \mathbb{R}_{t+1}^k + \tau_{t+1} \gamma \mathbb{R}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t + (1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \left(1 - (1 - \alpha_{t+1}^k)^2 \right) \right) q_{t+1}, \quad k = 1, \dots, t \quad (\Delta \mathbb{R}_t^k) \\
\mathbb{R}_t &= \mathbb{R}_{t+1} + (1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \right) q_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t, \quad (\Delta \mathbb{R}_t) \\
\mathbb{H}_t &= \left(1 - \frac{1}{t_h} \right) \mathbb{H}_{t+1} + \frac{1 - m_h}{t_h} \mathbb{R}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t. \quad (\Delta \mathbb{H}_t)
\end{aligned}$$

²⁷In fact, Equation ($FOC - \alpha_t^k$) can be rewritten as $\alpha_t^k = A_t^k + \frac{\kappa}{s_t^k + i_t^k + r_t^k} \left(\frac{1}{\alpha_t^k} - \frac{1}{1 - \alpha_t^k} \right)$ for a certain number A_t^k , which is independent of κ . The equation defines a polynomial of degree 3. Then, it is possible to show that the solution in trigonometric functions can be obtained using the Viète's formula, and that α_t^k converges to $\max\{0, \min\{1, A_t^k\}\}$ as $\kappa \downarrow 0$.

As discussed in Section 4.2, the system of adjoint equations is backwards with the following boundary condition:

$$\lim_{t \rightarrow \infty} \frac{\mathbb{S}_t^k}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{I}_t^k}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{R}_t^k}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{I}\mathbb{T}_t}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{R}\mathbb{T}_t}{p_t \delta_A^{t-1}} = \lim_{t \rightarrow \infty} \frac{\mathbb{H}_t}{p_t \delta_A^{t-1}} = \delta_A \left(1 + \frac{c}{2}\right) \forall k.$$

As usual, we write $\Delta Y_t := ((\Delta \mathbb{S}_t^k)_{k \leq t+1}, (\Delta \mathbb{I}_t^k)_{k \leq t}, \Delta \mathbb{I}\mathbb{T}_t, (\Delta \mathbb{R}_t^k)_{k \leq t}, \Delta \mathbb{R}\mathbb{T}_t, \Delta \mathbb{H}_t)$ for the system of adjoint equations. For the fixed aggregates $(S_t, I_t, \widehat{S}_t, \widehat{I}_t)$ and policies (λ_t, τ_t) , Equations $(\Delta \pi_t)$, (ΔY_t) and $(FOC - \alpha_t^k)$ determine a representative agent's best response, that is her rate of social distancing $(\alpha_t^k)_{k \leq t}$ as well as the individual state and adjoint vectors, π_t and Y_t , respectively. Since the agents are ex-ante identical, the Law of large numbers implies that in equilibrium the individual state variables match exactly the aggregate state variables, that is $\pi_t = \Pi_t$, in particular: $(s_t, i_t, \widehat{s}_t, \widehat{i}_t) = (S_t, I_t, \widehat{S}_t, \widehat{I}_t)$.

10.2 Derivation of the government's optimal policy

The government's optimization problem is described in details in Section 5: The government chooses the rates of lockdown, social distancing and testing as well as the aggregate state variables and agent's adjoint variables to maximize her own payoff. In this task the government is constrained by the fact every agent must indeed find it in her interest to follow the designated social distancing policy, that is the agent's optimality conditions must be respected. In addition, the total test count must be exactly equal to the capacity, i.e., \overline{X}_t . Thus, we arrive at the following problem:

$$\max_{\lambda_t, (\alpha_t^k)_{k \leq t}, \tau_t, \Pi_t, Y_t} Y_G + \xi \cdot (1 - D) \quad \text{subject to} \quad (\Delta \Pi_t), (\Delta Y_t), (FOC - \alpha_t^k) \text{ and } (F - \tau_t).$$

We now set up the Lagrangian for the problem denoting the undiscounted dual variables associated with Equations $(\Delta \Pi_t)$ by $\overline{Y}_t = ((\overline{\mathbb{S}}_t^k)_{k \leq t+1}, (\overline{\mathbb{I}}_t^k)_{k \leq t}, \overline{\mathbb{I}\mathbb{T}}_t, (\overline{\mathbb{R}}_t^k)_{k \leq t}, \overline{\mathbb{R}\mathbb{T}}_t, \overline{\mathbb{H}}_t)$ and those associated with Equations (ΔY_t) by $((\overline{\mathbb{S}}_t^k)_{k \leq t}, (\overline{\mathbb{I}}_t^k)_{k \leq t}, \overline{\mathbb{I}\mathbb{T}}_t, (\overline{\mathbb{R}}_t^k)_{k \leq t}, \overline{\mathbb{R}\mathbb{T}}_t, \overline{\mathbb{H}}_t)$. In addition, we let $(\eta_t^k)_{k \leq t}$ and χ_t be the Lagrange multipliers on

$(FOC - \alpha_t^k)$ and $(F - \tau_t)$, respectively. Taking all pieces together, the Lagrangian can be written as follows:

$$\begin{aligned}
& \underbrace{(1 - \delta_G) \sum_{t=1}^{\infty} \delta_G^{t-1} \left(Y_t q_t + \frac{\delta_G + \xi / \delta_G^{t-1}}{1 - \delta_G} (1 - D_{t+1}) p_t \right)}_{Y_G} + \underbrace{\sum_{t=1}^{\infty} \chi_t \left(-\bar{X}_t + \tau_t \left(\gamma (S_t - \beta_w \lambda_t^2 S_t I_t - \beta_s \widehat{S}_t \widehat{I}_t) + \gamma \left(R_t + \frac{1 - m_t}{l_t} I_t \right) + \left(1 - \frac{1}{l_t} \right) I_t + \beta_w \lambda_t^2 S_t I_t + \beta_s \widehat{S}_t \widehat{I}_t \right) \right)}_{F - \tau_t} \\
& + \underbrace{\sum_{t=1}^{\infty} \sum_{k \leq t} \eta_t^k \left((1 - \delta_A) c \delta_A^{t-1} \left((1 - \alpha_t^k) (S_t^k + I_t^k + R_t^k) + \kappa \left(\frac{1}{\alpha_t^k} - \frac{1}{1 - \alpha_t^k} \right) \right) q_t - (1 - \delta_A) \delta_A^{t-1} \beta_s \left(\phi^+ S_t^k \widehat{I}_t + \phi^- I_t^k \widehat{S}_t \right) - \Delta_t^k \beta_s S_t^k \widehat{I}_t \right)}_{FOC - \alpha_t^k} \\
& + \underbrace{\sum_{t=1}^{\infty} \sum_{k \leq t} \bar{S}_t^k \left(-S_{t+1}^k + (1 - \tau_t \gamma) \left(S_t^k - \beta_w \lambda_t^2 S_t^k I_t - \beta_s \alpha_t^k S_t^k \widehat{I}_t \right) \right)}_{\Delta S_t^k} + \underbrace{\sum_{t=1}^{\infty} \bar{S}_t^{t+1} \left(-S_{t+1}^{t+1} + \tau_t \gamma \left(S_t - \beta_w \lambda_t^2 S_t I_t - \beta_s \widehat{S}_t \widehat{I}_t \right) \right)}_{\Delta S_t^{t+1}} \\
& + \underbrace{\sum_{t=1}^{\infty} \sum_{k \leq t} \bar{I}_t^k \left(-I_{t+1}^k + (1 - \tau_t) \left(\left(1 - \frac{1}{l_t} \right) I_t^k + \beta_w \lambda_t^2 S_t^k I_t + \beta_s \alpha_t^k S_t^k \widehat{I}_t \right) \right)}_{\Delta I_t^k} + \underbrace{\sum_{t=1}^{\infty} \bar{I}_t^{t+1} \left(-I_{t+1}^{t+1} + \left(\left(1 - \frac{1}{l_t} \right) I_t + \beta_w \lambda_t^2 S_t I_t + \beta_s \widehat{S}_t \widehat{I}_t \right) \right)}_{\Delta I_t^{t+1}} \\
& + \underbrace{\sum_{t=1}^{\infty} \sum_{k \leq t} \bar{R}_t^k \left(-R_{t+1}^k + (1 - \tau_t \gamma) \left(R_t^k + \frac{1 - m_t}{l_t} I_t^k \right) \right)}_{\Delta R_t^k} + \underbrace{\sum_{t=1}^{\infty} \bar{R}_t^{t+1} \left(-R_{t+1}^{t+1} + \left(R_t + \frac{1 - m_t}{l_t} I_t + \frac{1 - m_t}{h_t} H_t + \tau_t \gamma \left(R_t + \frac{1 - m_t}{l_t} I_t \right) \right) \right)}_{\Delta R_t^{t+1}} \\
& + \underbrace{\sum_{t=1}^{\infty} \bar{H}_t \left(-H_{t+1} + \left(\left(1 - \frac{1}{h_t} \right) H_t + \frac{m_t}{l_t} (I_t + I_t^{t+1}) \right) \right)}_{\Delta H_t} + \underbrace{\sum_{t=1}^{\infty} \sum_{k \leq t+1} \bar{S}_t^k \left(-S_{t+1}^k + (1 - \tau_{t+1} \gamma) S_{t+1}^k + \tau_{t+1} \gamma S_{t+1}^{t+2} - \Delta_{t+1}^k \left(\beta_w \lambda_{t+1}^2 I_{t+1} + \beta_s \widehat{I}_{t+1} \alpha_{t+1}^k \right) + \delta_A^t \left(1 + \frac{c}{2} \right) p_t \right)}_{\dots} \\
& + \underbrace{(1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \left(1 - (1 - \alpha_{t+1}^k)^2 \right) - \phi^+ \left(\beta_w \lambda_{t+1}^2 I_{t+1} + \beta_s \alpha_{t+1}^k \widehat{I}_{t+1} \right) \right) q_{t+1}}_{\Delta S_t^k} \\
& + \underbrace{\sum_{t=1}^{\infty} \sum_{k \leq t+1} \bar{I}_t^k \left(-I_{t+1}^k + \left(\left(1 - \frac{1}{l_t} \right) \left((1 - \tau_{t+1}) I_{t+1}^k + \tau_{t+1} \mathbb{I} \mathbb{T}_{t+1} \right) + \frac{1 - m_t}{l_t} \left((1 - \tau_{t+1}) R_{t+1}^k + \tau_{t+1} \gamma \mathbb{R} \mathbb{T}_{t+1} \right) + \frac{m_t}{l_t} \mathbb{H}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t \right) \right)}_{\dots} \\
& + \underbrace{(1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \left(1 - (1 - \alpha_{t+1}^k)^2 \right) - \phi^- \left(\beta_w \lambda_{t+1}^2 S_{t+1} + \beta_s \alpha_{t+1}^k \widehat{S}_{t+1} \right) \right) q_{t+1}}_{\Delta \bar{I}_t^k} + \underbrace{\sum_{t=1}^{\infty} \bar{I}_t^{t+1} \left(-I_{t+1}^{t+1} + \left(\left(1 - \frac{1}{l_t} \right) \mathbb{I} \mathbb{T}_{t+1} + \frac{1 - m_t}{l_t} \mathbb{R} \mathbb{T}_{t+1} + \frac{m_t}{l_t} \mathbb{H}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t \right) \right)}_{\Delta \mathbb{I} \mathbb{T}_t} \\
& + \underbrace{\sum_{t=1}^{\infty} \sum_{k \leq t+1} \bar{R}_t^k \left(-R_{t+1}^k + \left((1 - \tau_{t+1}) R_{t+1}^k + \tau_{t+1} \gamma \mathbb{R} \mathbb{T}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t + (1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \left(1 - (1 - \alpha_{t+1}^k)^2 \right) \right) q_{t+1} \right) \right)}_{\Delta \mathbb{R} \mathbb{T}_t} \\
& + \underbrace{\sum_{t=1}^{\infty} \bar{R}_t^{t+1} \left(-R_{t+1}^{t+1} + \left(\mathbb{R} \mathbb{T}_{t+1} + (1 - \delta_A) \delta_A^t \left(\lambda_{t+1} + \frac{c}{2} \right) q_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t \right) \right)}_{\Delta \mathbb{R} \mathbb{T}_t} + \underbrace{\sum_{t=1}^{\infty} \bar{H}_t \left(-H_{t+1} + \left(\left(1 - \frac{1}{h_t} \right) \mathbb{H}_{t+1} + \frac{1 - m_t}{h_t} \mathbb{R} \mathbb{T}_{t+1} + \delta_A^t \left(1 + \frac{c}{2} \right) p_t \right) \right)}_{\Delta \mathbb{H}_t}
\end{aligned}$$

In the problem above the following short notations are used: $S_t = \sum_{k \leq t} S_t^k$, $I_t = \sum_{k \leq t} I_t^k$, $R_t = \sum_{k \leq t} R_t^k$, $\widehat{S}_t = \sum_{k \leq t} \alpha_t^k S_t^k$, $\widehat{I}_t = \sum_{k \leq t} \alpha_t^k I_t^k$, $\widehat{R}_t = \sum_{k \leq t} \alpha_t^k R_t^k$, $Y_t = \lambda_t (S_t + I_t + R_t + RT_t)$, $\bar{Y}_t = \lambda_t (\bar{S}_t + \bar{I}_t + \bar{R}_t + \bar{RT}_t)$ and $D_t = 1 - (S_t + I_t + IT_t + R_t + RT_t + H_t)$. In addition, we let $\Delta_t^k = (1 - \tau_t \gamma) S_t^k + \tau_t \gamma S_t^{t+1} - (1 - \tau_t) I_t^k - \tau_t \mathbb{I} \mathbb{T}_t$ and $\bar{\Delta}_t^k = (1 - \tau_t \gamma) \bar{S}_t^k + \tau_t \gamma \bar{S}_t^{t+1} - (1 - \tau_t) \bar{I}_t^k - \tau_t \bar{\mathbb{I}} \mathbb{T}_t$.

The optimal lockdown is a unique maximizer of Equation (OPT- λ_t), which is given by

$$\begin{aligned}
\lambda_t = \arg \max_{\lambda \in [\bar{\lambda}, 1]} & \left\{ \underbrace{\left((1 - \delta_G) \delta_G^{t-1} Y_t q_t + (1 - \delta_A) \delta_A^{t-1} \bar{Y}_t q_t \right)}_{=: a_t} \lambda - \right. \\
& \left. - \underbrace{\left(\beta_w \sum_{k \leq t} \left(\bar{\Delta}_t^k S_t^k I_t + \Delta_t^k \bar{S}_t^k I_t \right) + \beta_w \chi_t \tau_t (1 - \gamma) S_t I_t + \beta_w (1 - \delta_A) \delta_A^{t-1} \left(\phi^+ \bar{S}_t I_t + \phi^- S_t \bar{I}_t \right) q_t \right)}_{=: b_t} \lambda^2 \right\}. \tag{OPT- λ_t }
\end{aligned}$$

Denote by a_t and b_t the first and second coefficients in Equation (OPT- λ_t), respectively. Then, the optimal value of λ_t can be succinctly written as

$$\lambda_t = \begin{cases} 1 & \text{if } b_t \leq 0, \\ \min \left\{ 1, \max \left\{ 0, \bar{\lambda}, \frac{a_t}{2b_t} \right\} \right\} & \text{if } b_t > 0. \end{cases}$$

As for the other remaining control variables: The first-order conditions with respect to the testing rate τ_t and social distancing $(\alpha_t^k)_{k \leq t}$ pin down the Lagrange multipliers on the corresponding constraints, χ_t and $(\eta_t^k)_{k \leq t}$, respectively. The former is given by the following:

$$\begin{aligned} & \chi_t \left(\gamma S_t + \left(1 - \frac{1}{t_i}\right) I_t + \gamma \left(R_t + \frac{1 - m_i}{t_i} I_t \right) + (1 - \gamma) \left(\beta_w \lambda_t^2 S_t I_t + \beta_s \widehat{S}_t \widehat{I}_t \right) \right) = \\ & = \left(\sum_{k \leq t} \left(\gamma \left(\overline{S}_t^k - \overline{S}_t^{t+1} \right) S_t^k + \left(1 - \frac{1}{t_i}\right) \left(\overline{I}_t^k - \overline{I} \mathbb{T}_t \right) I_t^k + \gamma \left(\overline{R}_t^k - \overline{R} \mathbb{T}_t \right) \left(R_t^k + \left(1 - \frac{1 - m_i}{t_i}\right) I_t^k \right) \right) \right) + \\ & + \sum_{k \leq t} \left(\gamma \left(S_t^k - S_t^{t+1} \right) \overline{S}_t^k + \left(1 - \frac{1}{t_i}\right) \left(\mathbb{I}_t^k - \mathbb{I} \mathbb{T}_t \right) \overline{I}_t^k + \gamma \left(\mathbb{R}_t^k - \mathbb{R} \mathbb{T}_t \right) \left(R_t^k + \left(1 - \frac{1 - m_i}{t_i}\right) I_t^k \right) \right) + \\ & + \left(\overline{I}_t^k - \overline{I} \mathbb{T}_t - \gamma \left(\overline{S}_t^k - \overline{S}_t^{t+1} \right) \right) \left(\beta_w \lambda_t^2 S_t^k I_t + \beta_s \alpha_t^k + \eta_t^k \right) S_t^k \widehat{I}_t + \left(\mathbb{I}_t^k - \mathbb{I} \mathbb{T}_t - \gamma \left(S_t^k - S_t^{t+1} \right) \right) \left(\beta_w \lambda_t^2 \overline{S}_t^k I_t + \beta_s \left(\alpha_t^k + \eta_t^k \right) \overline{S}_t^k \widehat{I}_t \right). \end{aligned} \quad (\text{OPT-}\tau_t)$$

We note that Equation (OPT- τ_t) uniquely determines the Lagrange multiplier χ_t for the interior choice of $\tau_t \in (0, 1)$. Recall that the testing rate is determined by Equation (F - τ_t) in a way that exactly \overline{X}_t are conducted at date t . As a result, as long as \overline{X}_t is positive and not too large, the implied testing rate is indeed interior.

The first-order condition of the government's problem with respect to the rate of social distancing is as follows:

$$\begin{aligned} & \eta_t^k (1 - \delta_A) \delta_A^{t-1} c \left(S_t^k + I_t^k + R_t^k + \kappa \left(\frac{1}{(\alpha_t^k)^2} + \frac{1}{(1 - \alpha_t^k)^2} \right) \right) q_t + \sum_{m \leq t} \eta_t^m \left((1 - \delta_A) \delta_A^{t-1} \beta_s \left(\phi^+ S_t^m I_t^k + \phi^- I_t^m S_t^k \right) q_t - \Delta_t^m \beta_s S_t^m I_t^k \right) = \\ & = \chi_t \tau_t (1 - \gamma) \beta_s \left(S_t^k \widehat{I}_t + I_t^k \widehat{S}_t \right) - \beta_s \left(\overline{\Delta}_t^k S_t^k + \Delta_t^k \overline{S}_t^k \right) \widehat{I}_t - \beta_s \sum_{m \leq t} \left(\overline{\Delta}_t^m S_t^m + \Delta_t^m \overline{S}_t^m \right) \overline{I}_t^k + (1 - \delta_A) \delta_A^{t-1} c (1 - \alpha_t^k) \left(\overline{S}_t^k + \overline{I}_t^k + \overline{R}_t^k \right) q_t - \\ & - (1 - \delta_A) \delta_A^{t-1} \beta_s \left(\phi^+ \overline{S}_t^k \widehat{I}_t + \phi^- \overline{I}_t^k \widehat{S}_t + \phi^+ \widehat{S}_t I_t^k + \phi^- \widehat{I}_t S_t^k \right) q_t. \end{aligned} \quad (\text{OPT-}\alpha_t^k)$$

Equation (OPT- α_t^k) can be used to solve for the Lagrange variable $(\eta_t^k)_{k \leq t}$. According to the aforementioned equation, $(\eta_t^k)_{k \leq t}$ must be a solution to the linear system, i.e., $\mathcal{A}_t \vec{\eta}_t = \mathcal{B}_t$ for the certain $t \times t$ matrix \mathcal{A}_t and $t \times 1$ vector \mathcal{B}_t . Thus, as long as the matrix \mathcal{A}_t is invertible the solution is unique, and it is given by

$$\vec{\eta}_t = \mathcal{A}_t^{-1} \mathcal{B}_t.$$

Moving on, recall that the vector of agent's adjoints $\Upsilon_t = ((S_t^k)_{k \leq t+1}, (\overline{I}_t^k)_{k \leq t}, \mathbb{I} \mathbb{T}_t, (R_t^k)_{k \leq t}, \mathbb{R} \mathbb{T}_t, \mathbb{H}_t)$ is associated to the Lagrange variables $\overline{\Delta} \overline{\Pi}_t = ((\overline{S}_t^k)_{k \leq t}, (\overline{I}_t^k)_{k \leq t}, \overline{I} \mathbb{T}_t, (\overline{R}_t^k)_{k \leq t}, \overline{R} \mathbb{T}_t, \overline{H}_t)$. The first-order condition with respect to Υ_t specifies the forward system of equations (see (Δ $\overline{\Pi}_t$)) describing the dynamics of $\overline{\Pi}_t$. For completeness, we copy the system from the main text:

$$\overline{S}_{t+1}^k = (1 - \tau_t \gamma) \left(\overline{S}_t^k - \beta_w \lambda_t^2 \overline{S}_t^k I_t - \beta_s \left(\alpha_t^k \overline{S}_t^k - \eta_t^k S_t^k \right) \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta \overline{S}_t^k)$$

$$\overline{S}_{t+1}^{t+1} = \tau_t \gamma \left(\overline{S}_t - \beta_w \lambda_t^2 \overline{S}_t I_t - \beta_s \left(\widehat{S}_t - \sum_{k \leq t} \eta_t^k S_t^k \right) \widehat{I}_t \right), \quad (\Delta \overline{S}_t^{t+1})$$

$$\overline{I}_{t+1}^k = (1 - \tau_t) \left(\left(1 - \frac{1}{t_i}\right) \overline{I}_t^k + \beta_w \lambda_t^2 \overline{S}_t^k I_t + \beta_s \left(\alpha_t^k \overline{S}_t^k - \eta_t^k S_t^k \right) \widehat{I}_t \right) \text{ for } k = 1, \dots, t, \quad (\Delta \overline{I}_t^k)$$

$$\overline{I} \mathbb{T}_{t+1} = \left(1 - \frac{1}{t_i}\right) \overline{I} \mathbb{T}_t + \tau_t \left(\left(1 - \frac{1}{t_i}\right) \overline{I}_t + \beta_w \lambda_t^2 \overline{S}_t I_t + \beta_s \left(\widehat{S}_t - \sum_{k \leq t} \eta_t^k S_t^k \right) \widehat{I}_t \right), \quad (\Delta \overline{I} \mathbb{T}_t)$$

$$\overline{R}_{t+1}^k = (1 - \tau_t \gamma) \left(\overline{R}_t^k + \frac{1 - m_i}{t_i} \overline{I}_t^k \right), \text{ for } k = 1, \dots, t, \quad (\Delta \overline{R}_t^k)$$

$$\overline{R} \mathbb{T}_{t+1} = \overline{R} \mathbb{T}_t + \frac{1 - m_i}{t_i} \overline{I} \mathbb{T}_t + \frac{1 - m_h}{h_t} \overline{H}_t + \tau_t \gamma \left(\overline{R}_t + \frac{1 - m_i}{t_i} \overline{I}_t \right), \quad (\Delta \overline{R} \mathbb{T}_t)$$

$$\overline{H}_{t+1} = \left(1 - \frac{1}{t_h}\right) \overline{H}_t + \frac{m_i}{t_i} (\overline{I}_t + \overline{I} \mathbb{T}_t), \quad (\Delta \overline{H}_t)$$

and $\bar{I}_{t+1}^{t+1} = \bar{R}_{t+1}^{t+1}$ at all dates. The initial condition is that $\bar{\Pi}_t = 0$.

It remains only to identify the first-order condition with respect to the state variables. The vector of state variables $\Pi_t = ((S_t^k)_{k \leq t}, (I_t^k)_{k \leq t}, IT_t, (R_t^k)_{k \leq t}, RT_t, H_t)$ is associated to the vector of Lagrange multipliers \bar{Y}_t , which is $\bar{Y}_t = ((\bar{S}_t^k)_{k \leq t+1}, (\bar{I}_t^k)_{k \leq t}, \bar{IT}_t, (\bar{R}_t^k)_{k \leq t}, \bar{RT}_t, \bar{H}_t)$. Thus, the first-order conditions with respect to Π_t outputs the backward system of equations that describes the dynamics of \bar{Y}_t . The system is as follows:

$$\begin{aligned} \bar{S}_t^k &= \left((1 - \tau_{t+\gamma}) \bar{S}_{t+1}^k + \tau_{t+\gamma} \bar{S}_{t+1}^{t+2} \right) + \chi_{t+1} \tau_{t+1} \left(\gamma + (1 - \gamma) \left(\beta_w \lambda^2 I_{t+1} + \beta_s \widehat{I}_{t+1} \alpha_{t+1}^k \right) \right) + & (\bar{S}_t^k) \\ &+ (1 - \delta_G) \delta_G^t \lambda_{t+1} q_{t+1} + (\delta_G^t + \xi) p_t + \eta_{t+1}^k \left((1 - \delta_A) \delta_A^t c (1 - \alpha_{t+1}^k) q_{t+1} - (1 - \delta_A) \delta_A^t \beta_s \phi^+ \widehat{I}_{t+1} q_{t+1} - \Delta_{t+1}^k \beta_s \widehat{I}_{t+1} \right) - \\ &- \sum_{m \leq t+1} \eta_{t+1}^m (1 - \delta_A) \delta_A^t \beta_s \phi^- I_{t+1}^m \alpha_{t+1}^k - \bar{\Delta}_{t+1}^k \left(\beta_w \lambda_{t+1}^2 I_{t+1}^k + \beta \alpha_{t+1}^m \widehat{I}_{t+1} \alpha_{t+1}^k \right) - \\ &- (1 - \delta_A) \delta_A^t \sum_{m \leq t+1} \bar{I}_{t+1}^m \phi^- \left(\beta_w \lambda_{t+1}^2 + \beta_s \alpha_{t+1}^m \right) q_{t+1}, \\ \bar{I}_t^k &= \left(1 - \frac{1}{t_i} \right) \left((1 - \tau_{t+1}) \bar{I}_{t+1}^k + \tau_{t+1} \bar{IT}_{t+1} \right) + \frac{1 - m_i}{t_i} \left((1 - \tau_{t+1}) \bar{R}_{t+1}^k + \tau_{t+1} \gamma \bar{RT}_{t+1} \right) + \frac{m_i}{t_i} \bar{H}_{t+1} + & (\bar{I}_t^k) \\ &+ \chi_{t+1} \tau_{t+1} \left(\gamma \frac{1 - m_i}{t_i} + 1 - \frac{1}{t_i} + (1 - \gamma) \left(\beta_w \lambda^2 S_{t+1} + \beta_s \widehat{S}_{t+1} \alpha_{t+1}^k \right) \right) + (1 - \delta_G) \delta_G^t \lambda_{t+1} q_{t+1} + (\delta_G^t + \xi) p_t + \\ &+ \eta_{t+1}^k (1 - \delta_A) \delta_A^t \left(c (1 - \alpha_{t+1}^k) - \beta_s \phi^- \widehat{S}_{t+1} \right) q_{t+1} - \sum_{m \leq t+1} \eta_{t+1}^m (1 - \delta_A) \delta_A^t \beta_s \left(\phi^+ + \Delta_{t+1}^m \right) S_{t+1}^m \alpha_{t+1}^k - \\ &- \sum_{m \leq t+1} \Delta_{t+1}^m \left(\beta_w \lambda_{t+1}^2 \bar{S}_{t+1}^m + \beta \alpha_{t+1}^m \bar{S}_{t+1}^m \alpha_{t+1}^k \right) - \sum_{m \leq t+1} \bar{\Delta}_{t+1}^m \left(\beta_w \lambda_{t+1}^2 S_{t+1}^m + \beta \alpha_{t+1}^m S_{t+1}^m \alpha_{t+1}^k \right) - \\ &- (1 - \delta_A) \delta_A^t \sum_{m \leq t+1} \bar{S}_{t+1}^m \phi^+ \left(\beta_w \lambda_{t+1}^2 + \beta_s \alpha_{t+1}^m \right) q_{t+1}, \\ \bar{IT}_t &= \left(1 - \frac{1}{t_i} \right) \bar{IT}_{t+1} + \frac{1 - m_i}{t_i} \bar{R}_{t+1} + \frac{m_i}{t_i} \bar{H}_{t+1} + (\delta_G^t + \xi) p_t, & (\bar{IT}_t) \\ \bar{R}_t^k &= (1 - \tau_{t+1}) \bar{R}_{t+1}^k + \tau_{t+1} \gamma \bar{RT}_{t+1} + \chi_{t+1} \tau_{t+1} \gamma + \eta_{t+1}^k (1 - \delta_A) c \delta_A^t (1 - \alpha_{t+1}^k) q_{t+1} + (1 - \delta_G) \delta_G^t \lambda_{t+1} q_{t+1} + (\delta_G^t + \xi) p_t & (\bar{R}_t^k) \\ \bar{RT}_t &= \bar{RT}_{t+1} + (1 - \delta_G) \delta_G^t \lambda_{t+1} q_{t+1} + (\delta_G^t + \xi) p_t, & (\Delta \bar{RT}_t) \\ \bar{H}_t &= \left(1 - \frac{1}{t_h} \right) \bar{H}_{t+1} + \frac{1 - m_h}{t_h} \bar{RT}_{t+1} + (\delta_G^t + \xi) p_t. & (\Delta \bar{H}_t) \end{aligned}$$

As already standard, we write $\Delta \bar{Y}_t := ((\Delta \bar{S}_t^k)_{k \leq t+1}, (\Delta \bar{I}_t^k)_{k \leq t}, \Delta \bar{IT}_t, (\Delta \bar{R}_t^k)_{k \leq t}, \Delta \bar{RT}_t, \Delta \bar{H}_t)$ for the system of adjoint equations. We note that this system, as one for Y_t , is backwards with the following ‘‘initial’’ condition:

$$\lim_{t \rightarrow \infty} \frac{\bar{S}_t^k}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{I}_t^k}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{R}_t^k}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{IT}_t}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{RT}_t}{p_t} - \delta_G^t = \lim_{t \rightarrow \infty} \frac{\bar{H}_t}{p_t} - \delta_G^t = \xi \quad \forall k.$$

10.3 Calibration

In this section we discuss in details the rationale behind the assumed government policies and the set of fixed parameters which are used to calibrate the parameters of interest.

10.3.1 Lockdown function

We start with the approximation of actual lockdown (see Figure 3a) which is assumed to take the following parametric form:

$$\lambda_t = \begin{cases} 1 & t = 1, \dots, 16, \\ \sqrt{1 - \frac{1 - \bar{\lambda}}{\left(\frac{32}{17} - 1\right)^{\sigma_1}} \left(\frac{t}{55} - 1\right)^{\sigma_1}} & t = 17, \dots, 32, \\ \bar{\lambda} & t = 33, \dots, 103, \\ \sqrt{\bar{\lambda}^2 \sigma_2^{t-1-103} + (1 - \sigma_2)^{t-1-103}} & t = 104, \dots, T, \end{cases}$$

where $1 - \bar{\lambda}$ reflects the maximum level of lockdown that can be imposed by the government, while σ_1 and σ_2 controls its severity. The assumed lockdown function has four broad phases: (i) between Feb 29, 2020 ($t = 1$) and Mar 15, 2020 ($t = 16$), the US government did not impose any lockdown and the economy worked at full capacity, (ii) severe lockdown including bans on international travel were imposed between Mar 16 ($t = 17$) and Mar 31, 2020 ($t = 32$), reflected by the steep decline in the lockdown function.²⁸ The flat line (phase (iii)) between March 31 ($t = 32$) to June 5 ($t = 103$) reflects the severe lockdown imposed in majority of the states and cities. Moreland et al. [2020] provides a comprehensive analysis of lockdowns imposed by various states in the US. They show that about 80% of the states had mandatory stay-at-home orders in some form by March 31st, 2020. Most of these orders were gradually rescinded or allowed to exhaust from June 6 ($t = 103$) onwards. This is reflected in the gradual rise in phase (iv) of the assumed lockdown function.

The choice of σ_1 and σ_2 reflects the curvature of the lockdown policy. After the proclamation of national emergency, there was a flurry of measures taken to curb down international and interstate travel, as well as closure of most offices. This is denoted by a steep curvature in phase (ii), reflected by a high value of $\sigma_1 = 2$ ensuring approximately the maximum level of lockdown at the start of phase (iii). On the other hand, the relaxation of lockdown has been much slower, reflected in a gradual increase in the lockdown function at the fixed rate of σ_2 in phase (iv). We set $\sigma_2 = 0.9995$, so that 60% of the economy is functioning by Fall, 2021.

10.3.2 Test count

As explained above in Section 6, measuring the total number of tests and the number of positive and negative cases has been a tough task for researchers. The difficulty is driven by the fact that in the US there is no Federal standard for reporting testing data, and each individual state can decide to report a number of tests in terms of a) people tested at least once (unique people), b) samples collected and c) testing encounters. The "COVID Tracking Project" aggregates the state level data as follows: for each state, a total number of tests equals to a number of testing encounters if it is reported, otherwise it equals to a number of unique people if it is reported, otherwise it equals to a number of samples collected.

In order to calibrate and estimate the model, we aggregate the testing data by imputing a number of unique people reported by each state. For each state s , let $\bar{X}_{s,t}$ be the number of unique tested at time t if it is reported by this state; otherwise, let $\bar{X}_{s,t}$ be the total number of test in state s at time t multiplied by the average ratio of the total number of

²⁸This was the period in which the US government formally acknowledged Covid-19 to be a national emergency. [Link](#) to Proclamation 9994 by President Donald J. Trump.

tests to the number of unique people at time t across states which reported both. Then, we simple set

$$\bar{X}_t := \sum_{s \in \text{States}} \bar{X}_{s,t}.$$

The main reason for working with unique people is that this metric is reported by most states in our dataset, thus its imputation is more reliable; in contrast, testing encounters were used only by a few smaller states. We do however re-calibrate and re-estimate the model using the aggregated testing data provided by the "COVID Tracking Project". It turns out that both approaches yield very similar results. Instead of reporting all the estimation and optimal policy results again, we simply plot four time series of lockdown, social distancing, infections and deaths for these two specifications of deaths—see Figure 13.

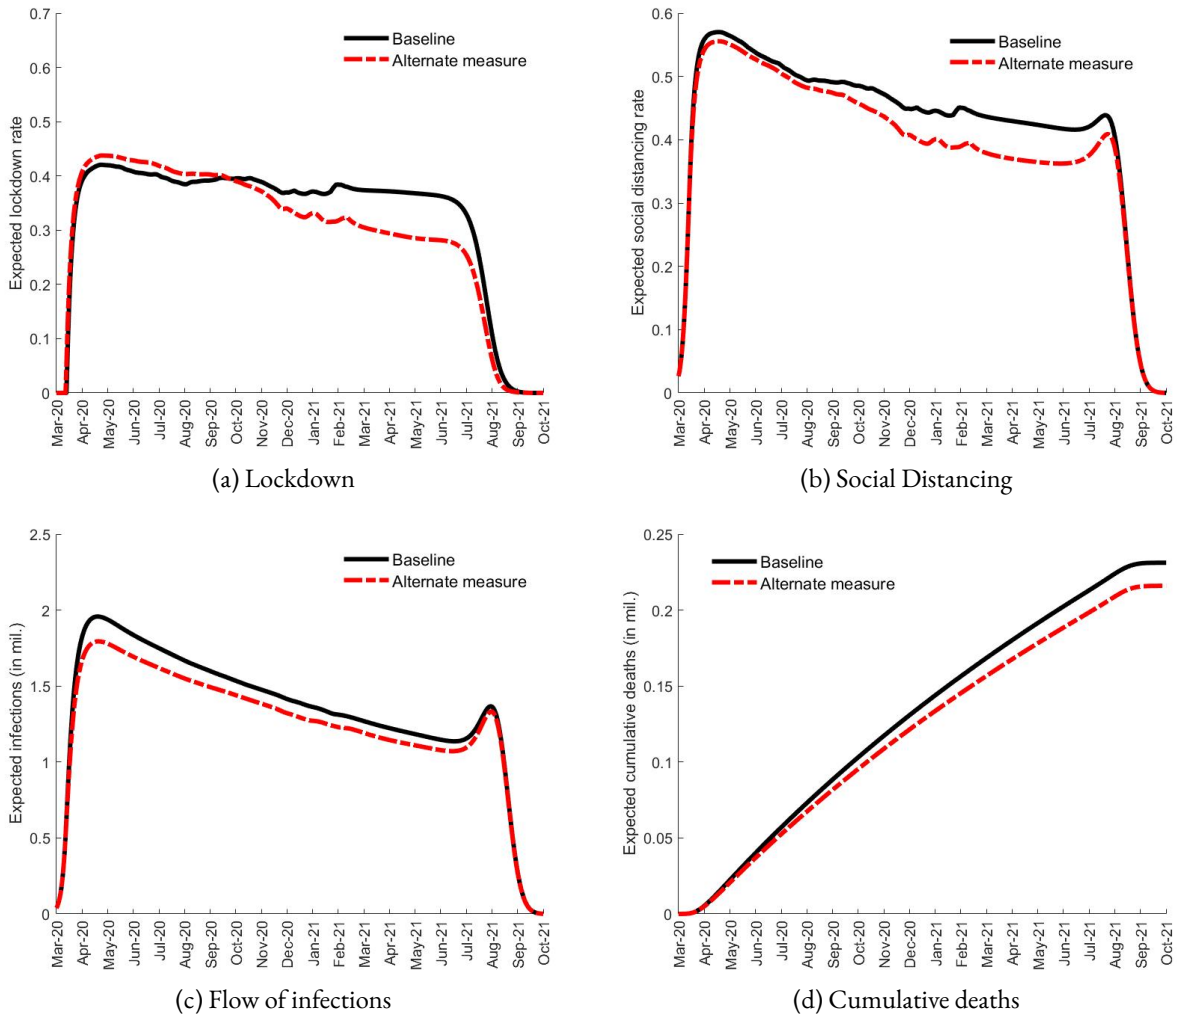


Figure 13: Optimal policy when the total number of tests is measured differently

The total number of tests available almost doubles in comparison to the baseline. However, in order to fit the data, the efficacy of tracing γ must become worse. So the gains in the total number of deaths in non-trivial but not large.

10.3.3 Fixed parameters

We now move to the fixed parameters shown in Table 2. First of all, we utilize the collective wisdom of various medical studies to calibrate the medical parameters. The average number of days of infection t_i is set to 10 days which includes the incubation and latency period. Conditional on developing further severe symptoms, the average number of days

of hospitalization (or at home critical care) is assumed to be $t_h = 7$. Our main basis for these numbers come from the [CDC Planning Scenario](#) dated March 19, 2021. For instance, their estimate for the mean time from exposure of the disease to the onset of symptoms is around 8 days. Similarly, the population weighted mean time between the onset of symptoms to hospitalization is around 9 days. The average hospitalization period is calculated similarly.

Since the time series of data on hospitalizations and daily deaths is available in the United States, and majority of Covid-19 related deaths have taken place at hospitals, we do a simple OLS regression between the two variables to calculate the rate at which those who develop severe symptoms are likely to die, which turns out to be $m_h = 0.1705$. This basically means that everything else being equal, 17% of all people who develop a severe case of Covid-19 will end up dying. Technically this parameter is not fixed because we estimate it from the data we have, but once estimated it will be fixed for the grand-estimation procedure.

We calculate the rate at which an infected person can develop severe symptoms and require extra care, that is the rate of transition from the state (**I**) to the state (**H**), to be $m_i = 0.0176$ using the following argument: First, the IFR or the infection fatality rate is taken to be 0.003. This means, ceteris paribus, there is a 0.3% chance of death upon contracting the virus. This corroborates with estimates obtained in some of the medical literature (see for instance [Ioannidis \[2021\]](#)). Note that the value of IFR is by no means settled in the medical literature. For instance, the CDC Planning Scenario of 19th March, 2021 pegs IFR to be around 0.004. [Alvarez et al. \[2021\]](#) uses a value of 0.0005, while [Farboodi et al. \[2021\]](#) uses an IFR of 0.0062. If anything, our choice of 0.003 may turn out to be on the higher side when the dust settles and one widely agreed upon number emerges. Once IFR is fixed, m_i is calculated using the identity $m_i \times m_H = IFR$.

The arrival of vaccine and more broadly vaccine dissemination that culminates in the date when the pandemic can be declared to have ended is modelled through a negative binomial distribution with mean 540 days and variance 180 days (see the entries for $\mathbb{E}T$ and $\mathbb{V}T$ in Table 2). This ensures that the pandemic will surely not end till around 404 days and it will almost certainly end in 593 days from the first day of the virus spread. This roughly squares with the prevailing wisdom, aggregated well in [New York Times \[2020\]](#), which had pegged the best case scenario for ending the pandemic to be the summer of 2021. Since the arrival of vaccine in our model coincides with the end of the pandemic, a mean of 540 days implies that the pandemic is expected to end in the Fall of 2021 (around August or September, 2021). At the current pace of vaccination in the US, it is expected that around 80% of the US population will be vaccinated by September, 2021 (see [New York Times \[2020\]](#)).

The maximum amount of lockdown the government can impose is capped at 70 percent. So $\bar{\lambda} = 0.3$ means that thirty percent of economic activities need to happen for the basic ingredients of society to keep functioning. We have in mind health, retail, government, utilities, and food manufacturing. This is the number used, for example, by [Alvarez et al. \[2021\]](#) and the uniform lockdown policy used in [Acemoglu et al. \[2021\]](#). Then, as standard in the literature, the government is forward-looking and discounts the future in a way that the annual interest rate equals to 5 percent. Our model allows for the flexibility of choosing any level of discounting by the agent, we calibrate the model for $\delta_A = \delta_G$.

Finally, we note that there is some indeterminacy in jointly identifying the total prevalence and the initial number of infected agent, see [Korolev \[2021\]](#). The standard approach in the literature is either to fix both or fix one of these parameters and calibrate the other. We decided to calibrate the total prevalence by taking $e_o = 0.007$ percent, i.e., there were 23000 infected people at the end of February (roughly one month since the first positive case was reported in the US). It must be noted that the choice of the initial seed varies greatly in the literature, depending on the time frame of the analysis. For instance, [Eichenbaum et al. \[2021a\]](#) works with an initial seed as high as 328,000 while [Farboodi et al. \[2021\]](#) work with an initial seed of roughly 17,000.

10.3.4 Fixing $\delta_A = 0$

The agents are modeled to be as forward looking as the government. [McAdams \[2021\]](#) notes that some studies in the literature fix $\delta_A = 0$, that is take the agents to be myopic optimizers, both for tractability and as thumb rule which approximates some reality. To assess robustness of our results and make them comparable to other studies, we fix $\delta_A = 0$, re-estimate the model and then solve for the optimal policy. All the other fixed parameters are kept the same as before.

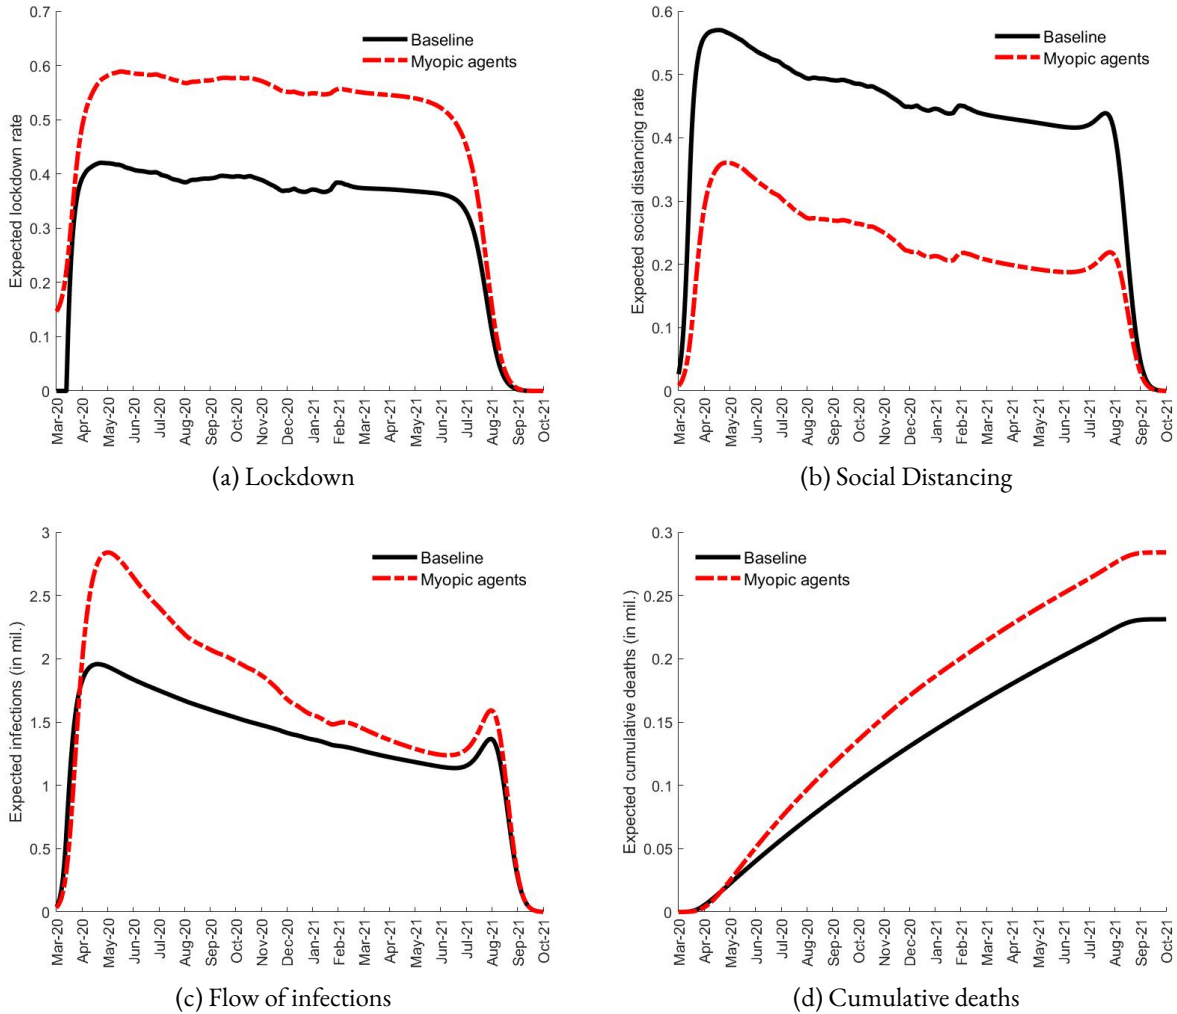


Figure 14: Results for myopic agents, $\delta_A = 0$

Since the agents are myopic, they internalize the costs of getting infected to a lesser extent and social distance less (Figure 14b). Expecting the reduced social distancing from the agents, the government thus locks down more (Figure 14a). Both social distancing and lockdown have the similar shape as the forward looking model — broadly persistent with a gradual decrease during the course of the pandemic, with a sharp decline towards the end. As a consequence of the agents' myopia, the number of infections and deaths is larger (Figures 14c and 14d). While the greater lockdown picks up some of the slack from reduced social distancing, the government still prefers not to completely lockdown at the maximal amount for a prolonged period for that would choke the economy. As before the optimal policy aims at reducing the effective reproductive number to one through the course of the pandemic.

References

- D. Acemoglu, V. Chernozhukov, I. Werning, and M. Whinston. Optimal targeted lockdowns in a multi-group sir model. *American Economic Review: Insights*, forthcoming, 2021.
- F. Alvarez, D. Argente, and F. Lippi. A simple planning problem for COVID-19 lock-down, testing, and tracing. *American Economic Review: Insights*, 3(3):367–82, 2021.
- A. Atkeson. A parsimonious behavioral SEIR model of the 2020 COVID epidemic in the United States and the United Kingdom. UCLA, 2021a.
- A. Atkeson. Behavior and the dynamics of epidemics. UCLA, 2021b.
- S. Bandyopadhyay, K. Chatterjee, K. Das, and J. Roy. Learning versus habit formation: Optimal timing of lockdown for disease containment. *Journal of Mathematical Economics*, 93(102452), 2021. doi: <https://doi.org/10.1016/j.jmateco.2020.11.008>.
- D. Baqaee, E. Farhi, M. Mina, and J. H. Stock. Policies for a second wave. *Brookings Papers on Economic Activity, Summer Special ed.*, 51, 2020.
- Z. Barnett-Howell and A. M. Mobarak. Should low-income countries impose the same social distancing guidelines as Europe and North America to halt the spread of COVID-19? Yale University, 2020.
- R. J. Barro, J. F. Ursúa, and J. Weng. The coronavirus and the great influenza pandemic: Lessons from the "Spanish Flu" for the coronavirus's potential effects on mortality and economic activity. NBER, 2020.
- D. Bergera, K. Herkenhoff, C. Huang, and S. Mongey. Testing and reopening in an SEIR model. *Review of Economic Dynamics*, 2021.
- A. Brodeur, D. Gray, A. Islam, and S. Bhuiyan. A literature review of the economics of COVID-19. *Journal of Economic Surveys*, 35(4):1007–1044, 2021.
- R. J. Caballero and A. Simsek. A model of endogenous risk intolerance and LSAPs: Asset prices and aggregate demand in a "COVID-19" shock. *The Review of Financial Studies*, 34(11):5522–5580, 2021.
- V. V. Chari, R. Kirpalani, and C. Phelan. The hammer and the scalpel: On the economics of indiscriminate versus targeted isolation policies during pandemics. *Review of Economic Dynamics*, 42:1–14, 2021.
- S. Correia, S. Luck, and E. Verner. Pandemics depress the economy, public health interventions do not: Evidence from the 1918 flu. The Federal Reserve and MIT, 2020.
- COVID-19 Mental Disorders Collaborators. Global prevalence and burden of depressive and anxiety disorders in 204 countries and territories in 2020 due to the COVID-19 pandemic. *The Lancet*, 398(10312):1700–1712, 2021. doi: [https://doi.org/10.1016/S0140-6736\(21\)02143-7](https://doi.org/10.1016/S0140-6736(21)02143-7).
- K. Dasaratha. Virus dynamics with behavioral responses. Harvard University, 2020.
- Deutsche Welle. Merkel: Coronavirus is Germany's greatest challenge since World War II. [Weblink](#), March 18, 2020.

- M. Eichenbaum, S. Rebelo, and M. Trabandt. The macroeconomics of epidemics. *The Review of Financial Studies*, 34(11):5149–5187, 2021a.
- M. S. Eichenbaum, S. Rebelo, and M. Trabandt. The macroeconomics of testing and quarantining. *NBER Working Paper No. 27104*, 2021b.
- S. Engle, J. Keppo, M. Kudlyak, E. Quercioli, L. Smith, and A. Wilson. The behavioral sir model, with applications to the swine flu and COVID-19 pandemics. University of Wisconsin-Madison, 2021.
- P. D. Fajgelbaum, A. Khandelwal, W. Kim, C. Mantovani, and E. Schaal. Optimal lockdown in a commuting network. *American Economic Review: Insights*, forthcoming, 2020.
- M. Farboodi, G. Jarosch, and R. Shimer. Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory*, forthcoming, 2021.
- E. P. Fenichel, C. Castillo-Chavez, M. G. Ceddia, G. Chowell, P. A. G. Parra, G. J. Hickling, G. Holloway, R. Horan, B. Morin, C. Perrings, M. Springborn, L. Velazquez, and C. Villalobos. Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences of the United States*, 108(15):6306–6311, 2011.
- J. Fernández-Villaverde and C. I. Jones. Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities. University of Pennsylvania and Stanford University, 2020.
- Financial Times. Covid-19 restrictions not affecting social distancing, says ONS. [Weblink](#), September 30, 2020.
- Financial Times. Sweden no longer stands out in pandemic, says architect of ‘no lockdown’ policy. [Weblink](#), November 12, 2021a.
- Financial Times. Vietnam abandons zero-Covid strategy after record drop in GDP. [Weblink](#), September 30, 2021b.
- E. Giannone, N. Paixao, and X. Pang. Pandemic in an interregional model– staggered restart. Penn State University and Bank of Canada, 2020.
- A. Glover, J. Heathcote, D. Krueger, and J.-V. Rios-Rull. Health versus wealth: On the distributional effects of controlling a pandemic. Federal Reserve Bank and University of Pennsylvania, 2021.
- V. Guerrieri, G. Lorenzoni, L. Straub, and I. Werning. Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? *American Economic Review*, forthcoming, 2021.
- T. Hale, N. Angrist, R. Goldszmidt, B. Kira, A. Petherick, T. Phillips, S. Webster, E. Cameron-Blake, L. Hallas, S. Majumdar, et al. A global panel database of pandemic policies (oxford covid-19 government response tracker). *Nature Human Behaviour*, 5(4):529–538, 2021.
- R. E. Hall, C. I. Jones, and P. J. Klenow. Trading off consumption and covid-19 deaths. Stanford University, 2020.
- J. P. Ioannidis. Infection fatality rate of covid-19 inferred from seroprevalence data. *Bulletin of the World Health Organization*, 99(1):19, 2021.
- C. Jones, T. Philippon, and V. Venkateswaran. Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies*, 34(11):5188–5223, 2021.

- G. Kaplan, B. Moll, and G. L. Violante. The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the u.s. University of Chicago, LSE and Princeton University, 2020.
- W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of London Series A*, 115(772):700–721, 1927.
- I. Korolev. Identification and estimation of the SEIRD epidemic model for COVID-19. *Journal of Econometrics*, 220(1):63–85, 2021.
- D. Krueger, H. Uhlig, and T. Xie. Macroeconomic dynamics and reallocation in an epidemic. University of Pennsylvania, University of Chicago and National University of Singapore, 2020.
- Mayo Clinic. COVID-19 and your mental health. [Weblink](#), November 21, 2021.
- D. McAdams. The blossoming of economic epidemiology. *Bulletin of the World Health Organization*, 13:539–570, 2021.
- A. Moreland, C. Herlihy, M. A. Tynan, G. Sunshine, R. F. McCord, C. Hilton, J. Poovey, A. K. Werner, C. D. Jones, E. B. Fulmer, et al. Timing of state and territorial covid-19 stay-at-home orders and changes in population movement—united states, march 1–may 31, 2020. *Morbidity and Mortality Weekly Report*, 69(35):1198, 2020.
- Y. Narita and A. Sudo. Curse of democracy: Evidence from the 21st century. Yale University, 2021.
- R. Neher, I. Aksamentov, N. Noll, J. Albert, and R. Dyrdak. Covid-19 scenarios. University of Basel, <https://covid19-scenarios.org/about>, 2020.
- New York Times. How long will a vaccine really take? [Weblink](#), April 30, 2020.
- J. Òscar., S. R. Singh, and A. M. Taylor. Longer-run economic consequences of pandemics. University of California, Davis, 2020.
- N. Perra, D. Balcan, B. Gonçalves, and A. Vespignani. Towards a characterization of behavior-disease models. *PLoS One*, 6(8):PMC3149628, 2011.
- A. Rizzo, M. Frasca, and M. Porfiri. Effect of individual behavior on epidemic spreading in activity-driven networks. *Physical Review E*, 90:042801, 2014.
- T. C. Schelling. The life you save may be your own. In S. B. Chase Jr., editor, *Problems in Public Expenditure Analysis*, pages 127–161. Brookings Institution, Washington D.C., 1968.
- The Gaurdian. Germany’s Covid-19 expert: ‘For many, I’m the evil guy crippling the economy’. [Weblink](#), April 26, 2020.
- F. Toxvaerd. Rational disinhibition and externalities in prevention. *International Economic Review*, 60(4):1737–1755, 2019.
- F. Toxvaerd. Equilibrium social distancing. University of Cambridge, 2020.