

INTERNAL AND EXTERNAL EFFECTS OF SOCIAL
DISTANCING IN A PANDEMIC

Maryam Farboodi
Gregor Jarosch
Robert Shimer

WORKING PAPER 27059

NBER WORKING PAPER SERIES

INTERNAL AND EXTERNAL EFFECTS OF SOCIAL DISTANCING IN A PANDEMIC

Maryam Farboodi
Gregor Jarosch
Robert Shimer

Working Paper 27059
<http://www.nber.org/papers/w27059>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020, Revised July 2020

We thank SafeGraph for making their data freely available to the research community. We are grateful for comments from Fernando Alvarez, Eric Budish, Luis Bettencourt, Charles I. Jones, Pete Klenow, Michael Kremer, Benjamin Moll, Simon Mongey, and Laura Pilossoph, as well as numerous seminar participants. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. Maryam Farboodi has received financial support from MIT for this research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Maryam Farboodi, Gregor Jarosch, and Robert Shimer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Internal and External Effects of Social Distancing in a Pandemic
Maryam Farboodi, Gregor Jarosch, and Robert Shimer
NBER Working Paper No. 27059
April 2020, Revised July 2020
JEL No. E1,H0,I1

ABSTRACT

We develop a quantitative framework for exploring how individuals trade off the utility benefit of social activity against the internal and external health risks that come with social interactions during a pandemic. We calibrate the model to external targets and then compare its predictions with daily data on social activity, fatalities, and the estimated effective reproduction number $R(t)$ from the COVID-19 pandemic in March-June 2020. While the laissez-faire equilibrium is consistent with much of the decline in social activity that we observed in US data, optimal policy further imposes immediate and highly persistent social distancing. Notably, neither equilibrium nor optimal social distancing is extremely restrictive, in the sense that the effective reproduction number never falls far below 1. The expected cost of COVID-19 in the US is substantial, \$12,700 in the laissez-faire equilibrium and \$8,100 per person under an optimal policy. Optimal policy generates this large welfare gain by shifting the composition of costs from fatalities to persistent social distancing.

Maryam Farboodi
MIT Sloan School of Management
100 Main Street, E62-627
Cambridge, MA 02142
and NBER
m.farboodi@gmail.com

Robert Shimer
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
shimer@uchicago.edu

Gregor Jarosch
Department of Economics
Princeton University
195 Julis Rabinowitz Building
Princeton, NJ 08544
and NBER
gregorjarosch@gmail.com

1 Introduction

A key parameter in workhorse models of disease transmission is the rate at which sick people infect susceptible people. During the COVID-19 outbreak in 2020, all policy measures against the disease were, in one way or another, aimed at reducing this rate. Our starting point is the observation that the disease is primarily transmitted through social interactions, which in turn depend on individuals' level of social activity, such as meeting with friends, dining at restaurants, shopping at retail stores, working in offices, and flying to travel destinations. These social activities always have costs and benefits, but the emergence of COVID-19 created a significant new cost, an elevated risk of disease transmission. We argue that standard economic tools are well-suited for modeling the behavioral response to COVID-19 and for assessing the costs and benefits of policies like widespread social distancing.

Our main contribution in this paper is to develop a simple model where individuals value social activity but face a risk of getting sick when they have social interactions with other people who are already infected. In a *laissez-faire* economy, forward-looking and rational individuals understand the risk of getting sick and the cost of the disease. This creates an internal benefit to social distancing, and individuals curtail their activity accordingly. Self-interested people do not, however, internalize the risk of making other people sick once they are already ill, an external benefit of social distancing. This creates a role for mandatory social distancing policies.

We evaluate a quantitative version of a model capturing these features and compare the predictions with data on social activity, disease dynamics, and fatalities in the US and other countries. We find that the internal benefits of social distancing alone reduced mortality by 70 percent in the US by July 1, 2020, from 1,925,000 to 564,000 deaths. A combination of altruistic behavior and mandatory widespread social distancing policies reduced it by an additional 77 percent, to the observed 129,000 deaths. On the other hand, a fully optimal policy would have resulted in only 5,500 deaths by that date, at the cost of a greater reduction in social activity.

We start the paper by using data from SafeGraph and Google to depict the evolution of social activity. Using several different metrics, we show that individuals in the US and elsewhere substantially reduced their social activity before state and local governments imposed the first “stay-at-home” or “shelter-in-place” restrictions.

Conversely, as the level of infection either passed its peak or remained relatively low, social activity increased before those restrictions were relaxed.¹ This is consistent with our hypothesis that individuals' desire to avoid getting sick is an important determinant of their social activity. Nevertheless, we also find that mandatory social distancing reduced activity below the level in laissez-faire settings. Our results are remarkably similar across metrics, data sets, and locations.

We then show how to integrate both individual behavior and policy analysis into an otherwise standard epidemiological model (Kermack and McKendrick, 1927). Living individuals are either susceptible to COVID-19, currently infected by it, or recovered. Susceptible individuals may catch the disease from infected individuals through the interaction between their social activity. We capture an important feature of COVID-19, that much of the transmission appears to be either pre-symptomatic or asymptomatic, through an assumption that individuals do not know when they are infected. Otherwise individuals are forward-looking and rational. They understand how their past behavior affects the probability that they are currently infected. They also understand how other people behave, the prevalence of the disease, and the risk of dying from it. In short, they know the costs and benefits of their own social activity. Finally, we assume that a cure which eliminates the health risk associated with the disease arrives stochastically, creating a strong incentive to delay getting sick.

We use optimal control theory to derive two ordinary differential equations (ODEs) which capture the internal cost of being susceptible and infected, along with an optimality condition which determines the level of social activity. Together with two standard differential equations from the epidemiological model, the resulting system of four ODEs fully summarize the model and can easily be solved on a computer.

We also derive the corresponding system of ODEs which characterize the symmetric Pareto optimal allocation. Optimal policy chooses a time path for the amount of social activity, recognizing the internal and external health consequences of social activity. A comparison with the laissez-faire benchmark elucidates that the internal effects of social distancing capture the fact that higher social activity increases the risk of susceptible people getting infected. The external effects capture that higher

¹As we write this version of the paper, infection rates are rising rapidly across much of the southern part of the US, including Florida, Texas, and Arizona. We expect that social activity will again decline in these parts of the country.

social activity increases the risk of infected people making susceptible people sick. While individuals weigh the risk of getting sick against the internal benefit of social activity, optimal policy also recognizes that purely self-interested individuals ignore the external cost of transmitting infections.

We then calibrate the model to US data. The calibration targets various epidemiological findings such as the initial growth rate of the disease, the duration of infectivity, and the fatality rate, along with the value of a statistical life. We use the quantitative model for two primary purposes. First, we highlight key observations about the laissez-faire economy and optimal policy which are robust to a wide range of values for the model parameters. Second, we quantify the cost of the pandemic and show how its level and decomposition into health and economic cost varies with policy choices.

Our most important observations are the following. First, the laissez-faire reduction in social activity due to the internalized cost of infection is strong. Compared to a naïve model without any social distancing, the peak outbreak is delayed and much lower, substantially reducing the expected number of fatalities. However, individuals only reduce activity once the risk of infection becomes non-negligible.

Second, and in contrast, optimal social distancing starts as soon as the disease emerges, discontinuously imposing social distancing, e.g. through stay-at-home orders. This discrete drop in activity delays the spread of infections and hence buys time. Because of the hope for a cure, this enormously reduces expected fatalities and yields a substantial welfare gain, albeit at the cost of a reduction in social activity.

Third, optimal social distancing is persistent, remaining in place for years or until a cure is found. This is the flip side of delaying the spread of infections, since it means that the share of people who are susceptible to the disease falls very slowly. Nevertheless, we find that social activity asymptotically returns to normal and infections stop as the population achieves herd immunity.

Fourth, social distancing is never too restrictive in the sense that the *effective reproduction number* $R(t)$ never falls far below 1. This is equivalent to saying that infections decline very slowly after they peak. Instead, optimal policy gradually allows social activity to climb back to normal and in doing so effectively rolls over infections. $R(t)$ behaves similarly in a laissez-faire setting.

The finding that $R(t)$ never falls far below 1 reflects a combination of two factors. First, the costs and benefits of social activity naturally stabilize the infection rate and the amount of social activity. If the infection rate is low, the risk of getting sick is small and so social activity is high. High levels of social activity then naturally lead to increased infections through epidemiological dynamics, which in turn reduce social activity through economic incentives. The opposite is true if the infection rate is high. When $R(t) = 1$, infections are constant by definition and so these forces are in balance. Second, since individuals do not know their disease status, social activity also depends on the relative likelihood of being susceptible versus infected. Because the infection rate is low in our calibrated model, the probability of being susceptible declines only slowly. That is, both state variables, the susceptible and infected rates, are slow-moving. Social activity then inherits this stickiness and only changes very slowly, which in turn implies that infections remain almost constant. This means that once the effective reproduction number $R(t)$ reaches 1, it stays close to it.

Using our calibrated model, we also derive quantitative results about the expected cost of COVID-19 and the benefits of policy. We first show that both the time path of the cost and its composition differ sharply between laissez-faire and optimal policy. Under laissez-faire, almost 90 percent of the costs associated with the pandemic are pure health costs, while optimal policy substantially reduces health costs but leads to much higher costs from social distancing. Further, under laissez-faire most of the costs are incurred during the first year of the pandemic. In contrast, optimal policy delays infections and hence has higher costs further in the future.

Overall, the expected cost of COVID-19 for the US is drastic. Under a laissez-faire path, the expected total cost amounts to 28 percent of one year's GDP. While large, this cost is still far below the cost of the no-response benchmark, highlighting that individual behavior is a powerful force in curtailing the fallout from the disease. Under the optimal policy, the expected total cost is significantly smaller but still substantial, 18 percent of one year's GDP. The expected discounted flow cost per capita under the optimal policy starts at \$13 per person per day and is still at \$4 per day two years later. In light of what is thus far known about the economic and health costs brought about by the pandemic, these numbers appear to be a plausible lower bound on the expected cost.

Finally, we validate the model by contrasting it with data along various dimensions. We first show that it accounts well for the overall decline in social activity which we measure in both the Google and SafeGraph data. The laissez-faire version of our model predicts a 40 percent decline in social activity occurring within a four week period. In the US, we observed a 30 percent decline in activity during the week of March 13, 2020, the first week with widespread public awareness of the dangers of the disease. Notably, there were few stay-at-home orders issued during that week, so we believe that this reflects the laissez-faire dynamics of the model. On the other hand, it would have been optimal to reduce social activity immediately by about 40 percent and maintain those restrictions for a much longer period of time.

Second, we contrast the aforementioned prediction that $R(t)$ stays close to 1 after peak infections with two independent measures of $R(t)$ which cover a wide variety of locations, Fernández-Villaverde and Jones (2020) and <https://rt.live>. We show that this prediction bears out remarkably well in the data. Following peak infection, the stock of infected individuals has been slowly rolled over in most locations, keeping $R(t)$ close to 1.

Third, we consider the evolution of fatalities in the US. We show that total US fatalities by late June were substantially below what would have occurred under laissez-faire. At the same time, the actual 129,000 deaths vastly exceed the 5,500 deaths that the US would have seen by July 1st if it had adopted the optimal policy starting on March 13th.

Overall, the models predictions align well with the data along various dimensions given what is thus far known about the disease. It is therefore a natural laboratory for quantitatively evaluating both policies that we have observed, like widespread social distancing, and policies that have been proposed but not yet implemented, including massive testing, test-trace-and-quarantine, and antibody testing.

The paper proceeds as follows: Section 2 relates our paper to the literature; Section 3 documents the evolution of social activity; Section 4 lays out the model and defines laissez-faire equilibrium as well as the optimum; Section 5 calibrates the model and presents our main quantitative results; Section 6 compares our model to data; Section 7 summarizes our main qualitative observations; Section 8 discusses the role of several key assumptions and modeling choices; and Section 9 concludes.

2 Related Work

Our basic approach builds on the susceptible-infected-recovered (SIR) model (Kermack and McKendrick, 1927). There exists a small body of pre-COVID-19 work which aims at integrating optimal control of economic or social activity with a model of disease dynamics (e.g. Philipson and Posner, 1993, 1995; Philipson, 2000; Kremer, 1996; Fenichel, 2013; Greenwood, Kircher, Santos and Tertilt, 2019; Toxvaerd, 2019). Prior to 2020, quantitative evaluations of epidemiological models with standard economic features have, to the best of our knowledge, been limited to HIV/AIDS, a very different disease than COVID-19.

There is now a rapidly growing body of work which applies a mix of the basic epidemiological model and dynamic optimization tools from economics to think about various aspects of COVID-19.² Relative to these papers, we focus on writing down the simplest possible model that can capture the costs and benefits of widespread social distancing in a manner that is quantitative and in line with available evidence. Our quantitative objective means that we focus throughout on the fact that the disease can be transmitted through many types of social interactions, only some of which are captured by consumption or labor supply. Fortunately, this allows us to map our model into newly available high-frequency data from Google and SafeGraph on social activity. This type of data only became available to researchers after the onset of COVID-19. Our focus on social activity also allows us to connect our model with the policy that was actually employed, social distancing.

Arguably what is most distinctive about our approach is that we offer a host of evidence on various aspects of both individual behavior and disease characteristics which we use to quantitatively evaluate the performance of the model. We discuss the evidence not only on social distancing but also on mortality and on the effective reproduction number $R(t)$. We also take advantage of the model's tractability to pro-

²An incomplete list includes Acemoglu, Chernozhukov, Werning and Whinston (2020), Alvarez, Argente and Lippi (2020), Atkeson (2020), Barro, Ursua and Weng (2020), Bethune and Korinek (2020), Bognanni, Hanley, Kolliner and Mitman (2020), Brotherhood, Kircher, Santos and Tertilt (2020), Budish (2020), Dewatripont, Goldman, Muraille and Platteau (2020), Eichenbaum, Rebelo and Trabandt (2020), Garibaldi, Moen and Pissarides (2020), Glover, Heathcote, Krueger and Rios-Rull (2020), Hall, Jones and Klenow (2020), Jones, Philippon and Venkateswaran (2020), Kaplan, Moll and Violante (2020), Keppo, Kudlyak, Quercioli, Smith and Wilson (2020), Krueger, Uhlig and Xie (2020), Piguillem and Shi (2020), Rowthorn and Toxvaerd (2020), and Toxvaerd (2020).

vide intuition for why the laissez-faire and optimal policies give rise to the outcomes that we observe.

Because of the significant overlap in the models, our paper leads to broadly similar qualitative conclusions as other papers in the rapidly growing literature on COVID-19. Individual optimality alone leads to social distancing, yet not enough. Optimal policy imposes immediate social distancing but allows the disease to spread until we reach herd immunity. Relative to these papers, we develop a unified and tractable approach to study both equilibrium and optimum, offer new evidence on equilibrium behavior in the early stages of the COVID-19 outbreak, and give a systematic assessment of the model's quantitative performance using a host of data on various aspects of the epidemic. Our quantitative prediction that optimal policy is persistent but not overly restrictive appears to be novel in this literature. In particular, we believe we are the first to connect these properties of optimal policy to the effective reproduction number $R(t)$, a fundamental concept in epidemiology, and show that it stabilizes close to 1.

At a more fine-grained level, our paper differs from most others in its information structure. We assume that newly-infected individuals do not know their infection status and so spread the disease unknowingly. This means that policy cannot discriminate between susceptible and infected people, which in turn means that social distancing policies must be wide-spread, rather than targeted on infected individuals. The aforementioned papers typically make different assumptions about individuals' information sets. On a related point, other papers focus on a broader set of policy tools. We deliberately focus on widespread social distancing because it was ubiquitous in the early stages of the COVID-19 outbreak and is likely to be proposed again for future outbreaks. Moreover, targeted policies like test-trace-and-quarantine still appear unavailable or dysfunctional many months after the outbreak.

The paper is also related to an older literature on social externalities, including Diamond and Maskin (1979) and Kremer and Morcom (1998). In particular, Diamond and Maskin (1979) introduce the distinction between a *quadratic* and a *linear matching technology*. With quadratic matching, additional social activity by others raises the likelihood of social contact and thus disease transmission for all individuals. For example, with more individuals in parks, restaurants, and public transit, any given trip to a park/restaurant/subway is more likely to lead to disease. Such a matching

function has a search externality that, traditionally, is viewed as positive (Diamond, 1982), but that turns negative in the context of disease. It stands in contrast to a linear search technology where an individual’s total social contacts merely depend on her own social activity and not on those of others. We believe that such a technology applies to cases where social contacts are explicitly sought out. We therefore argue that the quadratic technology is appropriate to model the dynamics of COVID-19, while a linear technology might be the right tool to model an epidemic like HIV.

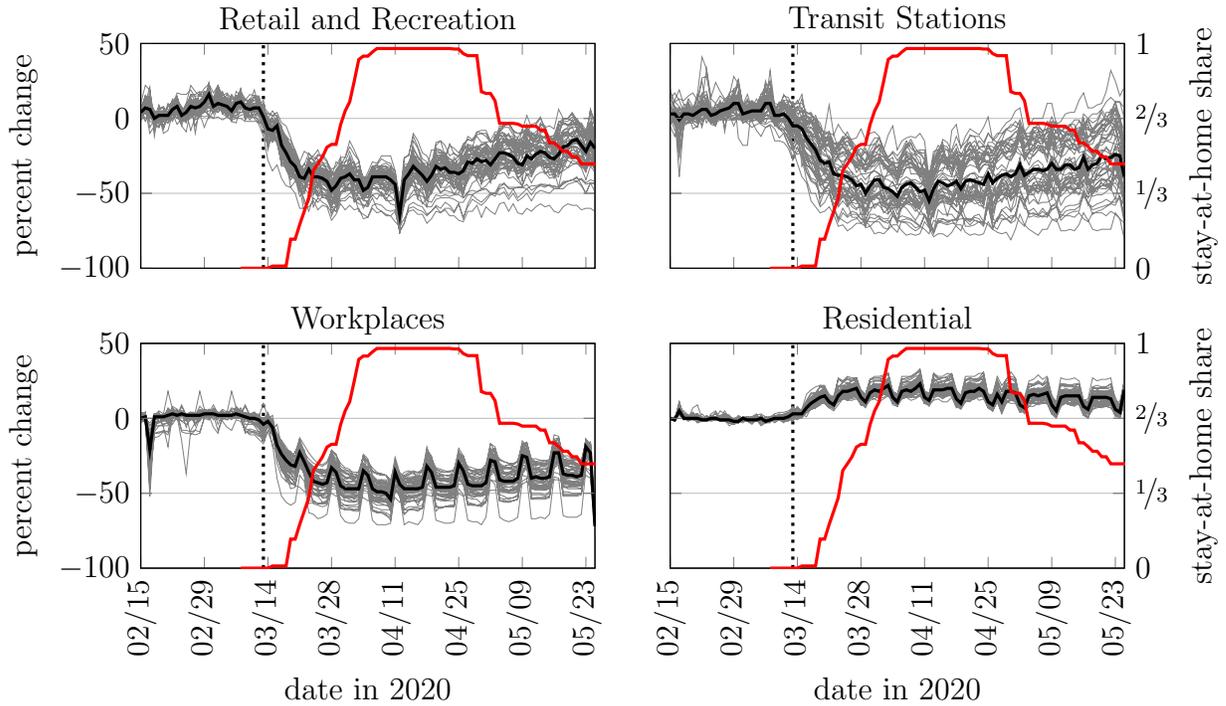
3 Declining Social Activity, Early and Everywhere

In this section, we use newly available data which document substantial social distancing across the US and in other countries even before any policy measures were taken. To do so, we use Google’s publicly available Community Mobility Reports complemented with data from SafeGraph. Both data sources are based on cellphone tracking data. Importantly, since we will model the choice of social activity directly, our setup can connect directly with this type of data.

Google Community Mobility Reports (<https://www.google.com/COVID19/mobility/>) use Google devices’ location history to track the number of visits and length of stay at places recognized as locations on Google maps. The activity metric is the percent change in activity relative to a baseline period of January 3 to February 6, 2020. We restrict attention to four categories of activity, each constructed by Google: “Retail and Recreation,” “Transit stations,” “Workplace,” and “Residential.”³ Figure 1 shows the percent change in activity in these categories across the US. The thin black lines in each figure plot Google’s activity metric across the fifty US states and the District of Columbia. The thick black line shows the cross-sectional median change in activity at any given date.

On March 13, 2020, indicated by a dotted line in Figure 1, the US government issued a proclamation declaring COVID-19 to be a national emergency. While this proclamation had little formal content, it reflected a growing awareness of the dangers

³We drop two categories, “Grocery and Pharmacy” and “Parks”. Both categories have complicated relationships with social distancing. For example, visits to grocery stores surged as people stocked their cupboards in the middle of March, while visits to parks depended on whether they were open in a particular location.



Notes: Each thin line corresponds to the Google Community Mobility Report for one of the 51 US states and the District of Columbia; see the text for a description of this measure. We retrieved this data from <https://www.google.com/COVID19/mobility/> on June 1, 2020. The thick black line marks the median state for that day. These are shown on the left hand scale. The red line shows the fraction of the US population subject to stay-at-home or shelter-in-place orders and is shown on the right hand scale. The population subject to these orders is based on authors' own calculations using <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html>. The vertical dotted line indicates March 13, 2020, the date of the declaration of national emergency.

Figure 1: Declining Activity, Early and Everywhere I

posed by the novel coronavirus in the US. In the following week, retail and recreation fell an average of 34 percentage points, transit stations by 27 percentage points, and workplaces by 29 percentage points. Over the same week, residential increased by 13 percent. All this occurred before any significant fraction of the US population was subject to stay-at-home orders, as indicated by the red lines in Figure 1. In the next week, as those orders spread across half the country, the first three categories fell by a further 10, 12, and 10 percentage points, respectively, while residential increased by another 5 percentage points. There was little change in the mobility index during the week after that.

After Easter Sunday, April 12, the trends started to reverse, with a noticeable increase in the first three mobility metrics and a decline in the residential metric. And as stay-at-home orders were eased in some states in late April and early May, there was no break in that trend. Instead, we saw a gradual ramp-up in the first three activities and a decline in the fourth during the second half of April and early May. We interpret this as confirming that there was widespread social distancing in the US, but that the majority of it was unrelated to mandatory stay-at-home orders. Instead, it reflected the desire of individuals to reduce their exposure to the disease.

Google Mobility Trends offers information on international locations as well. We use that to compare the evolution of social activity in Chicago with three Scandinavian capitals, Oslo, Copenhagen, and Stockholm. These three cities are arguably similar in a lot of ways, but they experienced quite different policies. While Norway and Denmark locked down on March 12 and 13, respectively, Sweden is widely considered an outlier in that it instituted few restrictions on social activity.

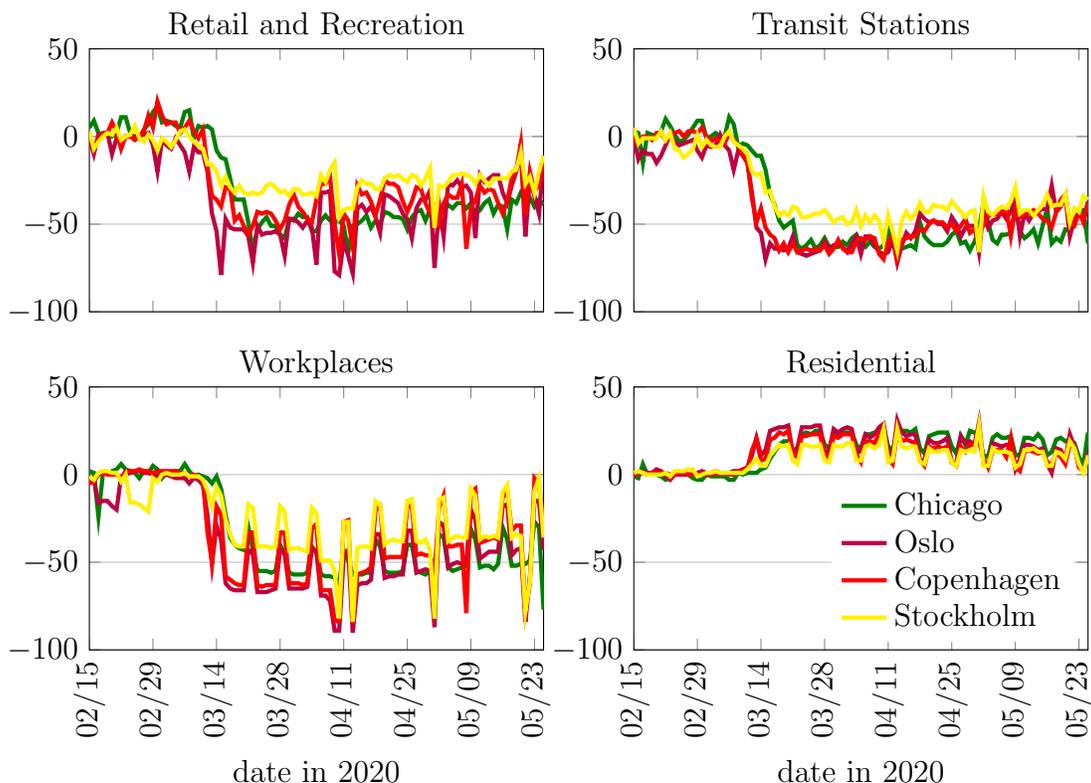
Figure 2 shows that social activity started falling in all four cities around March 13. The decline was sharpest in Oslo and Copenhagen, which were subject to formal stay-at-home orders, and more gradual in Stockholm and Chicago, which were not. Then on March 21, Illinois issued a shelter-in-place order. After that, the levels of retail, transit, and workplace activity in Chicago, Oslo, and Copenhagen were remarkably similar to each other and noticeably below Stockholm. Still, the significant decline in activity in Stockholm is consistent with our hypothesis that voluntary social distancing can substantially reduce activity even without formal stay-at-home orders.

We complement this analysis with data from SafeGraph.⁴ Among other things, SafeGraph provides disaggregated and detailed high-frequency information on individual travel in the US. The underlying data source is a panel of opt-in, anonymized smartphone devices, and is well balanced across US demographics and space.

In early April 2020, SafeGraph made several datasets freely available to researchers.⁵ The main dataset we use is their first “COVID-19 Response Dataset,” named “Weekly

⁴Attribution: SafeGraph, a data company that aggregates anonymized location data from numerous applications in order to provide insights about physical places. To enhance privacy, SafeGraph excludes census block group information if fewer than five devices visited an establishment in a month from a given census block group.

⁵For detailed information on the datasets we use, see <https://docs.safegraph.com/docs/weekly-patterns> and <https://docs.safegraph.com/docs/social-distancing-metrics>.

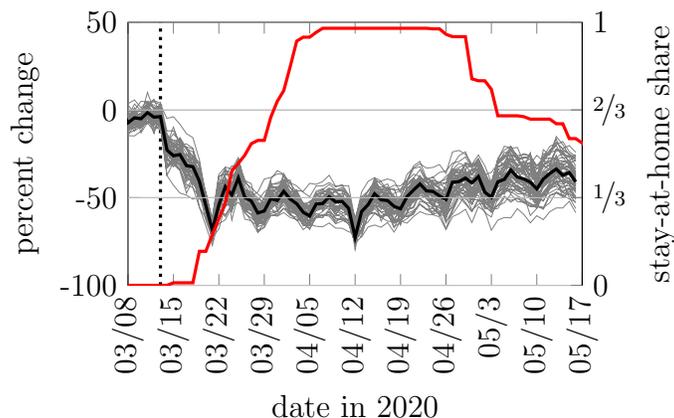


Notes: Figure constructed equivalently to Figure 1 but only these four cities. Activity metric is number of visits and length of stay at places recognized as locations on google maps, relative to a baseline period (January 3–February 6, 2020). Categories constructed by data provider. Retrieved from <https://www.google.com/COVID19/mobility/>.

Figure 2: Declining Activity, Early and Everywhere II

Patterns.” This registers GPS-identified visits to Points of Interest (POIs) (primarily businesses) with exact known location in the US at an hourly frequency in a balanced panel. We use data covering the time period from March 1 to May 16, 2020. The dataset is large: on March 1, the dataset recorded approximately 32.1 million individual visits to approximately 3.9 million POIs.

For each state and the District of Columbia, we count the total number of visits to POIs during each day. We express this variable as a percent change from a baseline week, the corresponding day during the first week of March. This gives us, for each state, a measure of the decline of social activity that naturally maps to the model. Figure 3 reports our results. We again plot thin lines for each state, as well as a thicker line for the median across states. The figure shows a remarkably uniform



Notes: Visits to Points of Interest in SafeGraph’s “Weekly Patterns” COVID-19 Response Dataset. We sum daily visits within each state and report the percent decline from the same day during the baseline week (first week of March 2020). Each thin line corresponds to one of the 51 US states and the District of Columbia. The thick black line marks the median state in terms of the decline relative to baseline at any given day. The solid red line is the percent of the population subject to stay-at-home or shelter-in-place orders, indicated on the right-hand axis. The population subject to these orders is based on authors’ own calculations using <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html>. The vertical dotted line indicates March 13, 2020, the date of the declaration of national emergency.

Figure 3: Declining Activity, Early and Everywhere III

contraction of social activity across US states beginning on March 13 and leveling off fifty percent below the baseline towards the end of March. Starting in late April, the data display a mild recovery that is somewhat more heterogeneous across states. These patterns are in line with the Google Community Mobility Reports.⁶

In Appendix A.1, we complement this with two other metrics from another dataset provided by SafeGraph, the fraction of devices that leave their home location at least once daily and the median time away from the home location, each at the census block level. These two variables display a somewhat smaller decline of 20 to 30 percent relative to baseline, perhaps natural given the difference in measures. The basic pattern, however, remains the same. Social activity started to contract substantially, rapidly, and before the first stay-at-home measures throughout the US.

⁶Figure 3 shows a sharp decline visits to POIs on March 20–22. For the median state, it falls from -41 on March 19 to -55 on March 20 to -69 on March 21 and then recovers equally abruptly to -54 on March 22 and -44 on March 23. This drop does not appear in the Google data and we are unsure what it represents. Ignoring those three days makes the similarity with Figure 1 even sharper.

One might be concerned that the decline in social activity reflects the closure of POIs rather than self-imposed restrictions on activity. To address this, we use the data to construct a measure of POI closures and then show that foot traffic declined substantially prior to closure; see Appendix A.2.

Finally, we note that in a recent working paper, Goolsbee and Syverson (2020) use SafeGraph data and conclude that, while social activity fell by 60 percentage points, stay-at-home orders explain only 7 percentage points of that. This is very much in line with our analysis here.

The bottom line is that we have direct and readily available measures of social activity. This is both the key endogenous variable in our model and the key driver of the pandemic. Since we model social activity rather than, say, consumption, our model can directly connect with this high-frequency data. In Section 6 we will confront the quantitative properties of our model with this evidence and argue that it offers a close account of the decline in social activity in March 2020.

4 Model

4.1 Setup

The basic epidemiological framework is a continuous time SIR model with a possibility of death, i.e. a SIRD model. We embed this into an economic model where individuals get utility from their social activity. Those who are susceptible to the disease can become infected through social interactions with infected individuals. The infected stochastically recover or die. Recovery confers immunity to the disease. Individuals do not know if they are susceptible or infected, but they can tell when they have recovered from the disease. Individuals are otherwise homogeneous.

Preferences over Social Activity Individuals discount the future at rate ρ and live forever unless they die from the disease. At any time $t \geq 0$, each living individual gets utility $u(a(t))$ from a level of *social activity* $a(t) \geq 0$. We assume u is single-peaked and its maximum is attained at a finite value a^* . Without loss of generality, normalize $a^* = 1$ and $u(a^*) = 0$. The first normalization keeps the notation the same as the basic SIR model in the absence of a behavioral response to a disease outbreak.

The latter allows us to use u as a measure of the utility cost of social distancing, i.e. of setting $a < a^*$.

Social Interactions and Disease Transmission At each point in time, individuals are in one of the four states, susceptible (s), infected (i), recovered (r), or deceased (d). Let $N_\theta(t)$, $\theta \in \{s, i, r, d\}$, denote the measure of individuals in each state. To make the environment nontrivial, assume that $N_i(0) > 0$, so there is a seed of infection, and $N_s(0) > 0$, so there are people who can get infected.

Disease transmission depends on *social interactions*, defined as the interaction between individuals' social activity. Let $A_\theta(t)$ denote the amount of social activity by each individual of type $\theta \in \{s, i, r\}$. We assume throughout that all individuals of a given type choose the same level of social activity, although when we study equilibrium, we consider the deviation of a single individual to a different level of activity. The rate at which some individual of type θ has a social interaction with some other individual of type θ' is proportional to $A_\theta(t)N_\theta(t)A_{\theta'}(t)N_{\theta'}(t)$, the product of the total amount of social activity of the two types. In particular, the rate at which susceptible individuals get infected is $\beta A_s(t)N_s(t)A_i(t)N_i(t)$, where $\beta > 0$ captures the ease of transmitting the disease. Formally, this is a quadratic matching technology with random search.⁷ We stress that there are not only social interactions between the susceptible and the infected, but also between all other groups. For instance, the recovered can optimally choose $A_r(t)$ without fear of infection; and susceptible individuals meet one another without consequences. Our model allows these interactions to happen, but implies that they do not affect the number of interactions where the disease gets transmitted.

The assumptions that preferences u depend on social *activity* while disease trans-

⁷An alternative would be a linear technology, in which case type θ individuals interact with type θ' individuals at rate $\frac{A_\theta(t)N_\theta(t)A_{\theta'}(t)N_{\theta'}(t)}{\sum_{\theta''} A_{\theta''}(t)N_{\theta''}(t)}$. This formulation makes sense if type θ individuals desire to interact with somebody at rate $A_\theta(t)$. With a linear technology, changing social activity by other people changes the distribution of social interactions without changing the level. Such an alternative might be the appropriate modeling choice for other epidemics such as HIV (Kremer and Morcom, 1998). See also Fenichel (2013). However, for COVID-19 it seems unlikely that additional social activity by non-infected individuals would reduce the infection risk of the susceptible, all else equal. We highlight, however, that this assumption is crucial for certain outcomes. For example, with a linear technology, it is optimal to boost the social activity of the recovered in order to moderate meetings between susceptible and infected people.

mission depends on social *interactions* are central to our view of social distancing. The former captures the idea that individuals value social activity (going to a restaurant, going for a walk, going to the office) and, absent health issues, are indifferent about whether other people are also engaging in social activity.⁸ On the other hand, if an individual goes for a walk and doesn't encounter anybody, they cannot get sick. Thus interactions are critical for disease transmission.

This disease transmission function captures the idea that if a particular group chooses little social activity, so $A_\theta(t)$ is small, then one is unlikely to make social contact with them. As a consequence, social activity displays a negative externality: it increases other people's social interactions, putting them at a higher risk of infection.

Recovery and Death Infected individuals recover at rate $(1 - \pi)\gamma$ and die at rate $\pi\gamma$, where $\pi \in [0, 1]$ is the infection fatality rate,⁹ the fraction of infected individuals who eventually die from the disease, and $1/\gamma$ is the expected duration of infection. We assume that at the moment an individual dies, he pays a utility cost v , the value of a statistical life. Equivalently, all individuals pay a utility cost $\kappa \equiv \pi v$ when they exit the infected state.

We assume that recovering from the disease confers lifetime immunity, so a recovered individual no longer transmits the disease and can no longer become sick. We recognize that the empirical evidence on the latter point is still unclear.

The End of the Disease We assume a cure for the disease is found at rate δ . For simplicity we assume that a cure, once found, is perfect and immediately wipes out all costs of the disease.

Information Individuals do not know whether they are susceptible or infected, but they know when they have recovered from the disease. We capture this through the measurability restriction that $A_s(t) = A_i(t)$ and use $A(t)$ to denote this common level

⁸For an environment where interactions are critical, see Diamond (1982). One can imagine reasons why the marginal utility of social activity is increasing or decreasing in the social activity of others.

⁹We assume that the infection fatality rate π is independent of the number of infected people, thus shutting down the possibility that the disease overwhelms the healthcare system. This is not meant to dispute that notion that “healthcare capacity” (or “ICU”) constraints are potentially an important feature of the COVID-19 outbreak. Instead, we do this to highlight that key properties—e.g. that optimal policy “flattens the curve”—arise even without such a constraint. We discuss this point further in Section 8.

of social activity. In Appendix C.2, we show that our results are quantitatively very similar if we assume that individuals also do not know when they have recovered from the disease, so $A_s(t) = A_i(t) = A_r(t)$.

Reproduction Number Our model has notions of two important epidemiological concepts, the basic and effective reproduction numbers. The *basic reproduction number* R_0 is defined as the expected number of susceptible people infected by a sick person in a world where almost everyone is susceptible and unaware of the disease. This was relevant, for example, when the disease first emerged. This is simply $R_0 = \beta/\gamma$. Throughout, we impose the natural parameter restriction that $\beta > \gamma$ since otherwise the disease never breaks out.

The *effective reproduction number* is the expected number of others infected by a sick person, given the current level of social activity $A(t)$ and the current fraction of susceptible people $N_s(t)$. This is $R(t) = R_0 A(t)^2 N_s(t)$. Social distancing drives the ratio of the effective reproduction number to the fraction of susceptible people, $R(t)/N_s(t)$, below the basic reproduction number R_0 . We show below that the number of infections grows (shrinks) if $R(t)$ is bigger (smaller) than 1.

4.2 Laissez-faire Equilibrium

In this section, we consider the problem of an individual choosing his own social activity, taking the social activity of others and the number of others with each health status as given. We then impose the equilibrium restriction that individual outcomes must coincide with the aggregate.

An individual has rational beliefs about his own probabilities of being susceptible, infected, and recovered, which we denote by $n_s(t)$, $n_i(t)$, and $n_r(t)$, respectively. The individual knows when he is recovered but cannot distinguish between the susceptible and infected states. He thus chooses two time paths for social activity, $a(t)$ when he is susceptible or infected and $a_r(t)$ when he is recovered. Finally, the individual discounts future utility at rate ρ and recognizes that the problem ends at rate δ when a cure is found. Putting this together, the individual solves

$$\max_{\{a(t), a_r(t)\}} \int_0^{\infty} e^{-(\rho+\delta)t} ((n_s(t) + n_i(t))u(a(t)) + n_r(t)u(a_r(t)) - \gamma n_i(t)\kappa) dt$$

subject to
$$n'_s(t) = -\beta a(t)n_s(t)A(t)N_i(t), \quad (1)$$

$$n'_i(t) = \beta a(t)n_s(t)A(t)N_i(t) - \gamma n_i(t), \quad (2)$$

$$n'_r(t) = (1 - \pi)\gamma n_i(t),$$

$$n'_d(t) = \pi\gamma n_i(t),$$

with $n_s(0) = N_s(0)$, $n_i(0) = N_i(0)$, $n_r(0) = N_r(0)$, and $n_d(0) = N_d(0)$ given. Note that through the choice of $a(t)$, an individual affects his own transitions across different health statuses. An individual transitions from susceptible to infected at a rate proportional to the product of his own activity level $a(t)$, his own probability of being susceptible $n_s(t)$, and the total activity by infected people $A(t)N_i(t)$. He exits the infected state at rate γ , either recovering or dying, paying an expected cost κ . The individual solves this problem taking the time path of $A(t)N_i(t)$ as given.

To solve the individual's problem, first note that $a_r(t)$ affects the objective but none of the constraints. It is thus optimal to set $a_r(t) = a^* = 1$ and thus $u(a_r(t)) = 0$ for all t . Dropping this control variable and the unnecessary third and fourth constraints, write the current value Hamiltonian as

$$\begin{aligned} H(n_s(t), n_i(t), a(t), \lambda_s(t), \lambda_i(t)) &= (n_s(t) + n_i(t))u(a(t)) - \gamma n_i(t)\kappa \\ &\quad - \lambda_s(t)\beta a(t)n_s(t)A(t)N_i(t) + \lambda_i(t)(\beta a(t)n_s(t)A(t)N_i(t) - \gamma n_i(t)), \end{aligned}$$

where $\lambda_s(t)$ and $\lambda_i(t)$ are the co-states associated with the two remaining constraints.

There are three necessary first order conditions for optimal control. First, the derivative of the Hamiltonian with respect to the control variable $a(t)$ is zero:

$$(n_s(t) + n_i(t))u'(a(t)) = (\lambda_s(t) - \lambda_i(t))\beta n_s(t)A(t)N_i(t). \quad (3)$$

This states that the returns from social activity balance the risk of getting infected. Additionally, the derivatives with respect to the state variables $n_s(t)$ and $n_i(t)$ are equal to minus the time derivative of the costate, with a correction for discounting:

$$(\rho + \delta)\lambda_s(t) - \lambda'_s(t) = u(a(t)) + (\lambda_i(t) - \lambda_s(t))\beta a(t)A(t)N_i(t), \quad (4)$$

$$(\rho + \delta)\lambda_i(t) - \lambda'_i(t) = u(a(t)) - \gamma(\kappa + \lambda_i(t)). \quad (5)$$

There are two more necessary conditions for optimality, the transversality conditions

$$\lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_s(t) n_s(t) = \lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_i(t) n_i(t) = 0. \quad (6)$$

Equilibrium requires that individual and aggregate behaviors are consistent at every point in time, $n_s(t) = N_s(t)$, $n_i(t) = N_i(t)$, and $a(t) = A(t)$ for all $t > 0$. Imposing those restrictions on equations (1)–(6) gives us

$$N'_s(t) = -\beta A(t)^2 N_s(t) N_i(t), \quad (7)$$

$$N'_i(t) = \beta A(t)^2 N_s(t) N_i(t) - \gamma N_i(t), \quad (8)$$

$$(N_s(t) + N_i(t)) u'(A(t)) = (\lambda_s(t) - \lambda_i(t)) \beta A(t) N_s(t) N_i(t), \quad (9)$$

$$(\rho + \delta) \lambda_s(t) - \lambda'_s(t) = u(A(t)) + (\lambda_i(t) - \lambda_s(t)) \beta A(t)^2 N_i(t), \quad (10)$$

$$(\rho + \delta) \lambda_i(t) - \lambda'_i(t) = u(A(t)) - \gamma (\kappa + \lambda_i(t)), \quad (11)$$

$$\lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_s(t) N_s(t) = \lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_i(t) N_i(t) = 0. \quad (12)$$

These equations and initial conditions for $N_s(0)$ and $N_i(0)$ fully summarize the model.

Note from equation (8) that the growth rate of infections, $N'_i(t)/N_i(t)$, is $\gamma(R(t) - 1)$, where $R(t) \equiv R_0 A(t)^2 N_s(t)$ is the effective reproduction number. Thus infections grow if and only if the effective reproduction number exceeds 1. This is why a low effective reproduction number is critical for reducing the prevalence of the disease.

We solve this model through a backward shooting algorithm. We provide a detailed description of our numerical solution algorithm in Appendix B

4.3 Social Optimum

We now solve the problem faced by a benevolent social planner who dictates the time path of social activity $A(t)$ and $A_r(t)$. The planner, like the individual, recognizes that a reduction in contacts lowers utility directly, but she also recognizes the externalities associated with infections. The planner solves

$$\max_{\{A(t), A_r(t)\}} \int_0^\infty e^{-(\rho+\delta)t} ((N_s(t) + N_i(t)) u(A(t)) + N_r(t) u(A_r(t)) - \gamma N_i(t) \kappa) dt \quad (13)$$

subject to equations (7) and (8) and $N_r'(t) = (1 - \pi)\gamma N_i(t)$. As in equilibrium, it is optimal to set $A_r(t) = a^* = 1$ and so $u(A_r(t)) = 0$ for all t , since this does not affect the evolution of the state variables. We can again write down a Hamiltonian. The necessary first order condition with respect to the control A is

$$(N_s(t) + N_i(t))u'(A(t)) = 2(\mu_s(t) - \mu_i(t))\beta A(t)N_i(t)N_s(t), \quad (14)$$

while the necessary costate equations are

$$(\rho + \delta)\mu_s(t) - \mu_s'(t) = u(A(t)) + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_i(t), \quad (15)$$

$$(\rho + \delta)\mu_i(t) - \mu_i'(t) = u(A(t)) + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_s(t) - \gamma(\kappa + \mu_i(t)). \quad (16)$$

Finally, the planner also has necessary transversality conditions

$$\lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \mu_s(t) N_s(t) = \lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \mu_i(t) N_i(t) = 0. \quad (17)$$

Together with the state equations (7) and (8), we again have a system of four differential equations and one static equation. We again solve it using a backward shooting algorithm, searching for the terminal value of $N_s(T)$ for given $N_i(T)$.

There are two key differences between the equilibrium equations (9)–(11) and the corresponding optimal equations (14)–(16). First, the planner recognizes that raising $A(t)$ increases meetings at rate proportional to $2A(t)$, while in equilibrium raising $a(t)$ increases meetings at rate proportional to $A(t)$. This creates an extra factor of 2 in equation (14) compared to equation (9). Second, the planner recognizes that sick people get other people sick, while in equilibrium individuals do not care about this outcome. This shows up as an extra term on the right hand side of equation (16) compared to equation (11).

4.4 Perfect and Imperfect Altruism

In the laissez-faire setting, people only care about their own health. In reality, diseases are often transmitted to family and friends, and so it seems plausible that many people would like to reduce the risk of transmitting the disease, not just the chance of getting it. We show how our basic setup can be extended to a model of imperfect altruism,

indexed by an altruism parameter $\alpha \in [0, 1]$. At one extreme, $\alpha = 0$, individuals would not socially distance if they knew they were sick. This is the laissez-faire equilibrium. At the other extreme, $\alpha = 1$, individuals fully internalize the cost of making others sick.

We modify the individual objective function to incorporate that individuals are concerned about making others sick:

$$\begin{aligned} \max_{\{a(t), a_r(t)\}} \int_0^\infty e^{-(\rho+\delta)t} & \left((n_s(t) + n_i(t))u(a(t)) + n_r(t)u(a_r(t)) \right. \\ & \left. - \gamma n_i(t)\kappa - \alpha\beta a(t)n_i(t)A(t)N_s(t)(\lambda_s(t) - \lambda_i(t)) \right) dt. \end{aligned}$$

The new piece is the last term. When an individual infects a susceptible person, at rate $\beta a(t)n_i(t)A(t)N_s(t)$, she suffers a utility loss equal to a fraction α of the difference $\lambda_s(t) - \lambda_i(t)$, where again $\lambda_\theta(t)$ is the costate variable on $n_\theta(t)$, $\theta \in \{s, i\}$. This difference represents the private cost of getting sick.

With this modification to the objective function, we can again write down the Hamiltonian and find the optimality and costate equations as well as the transversality conditions. Imposing $n_s(t) = N_s(t)$, $n_i(t) = N_i(t)$, and $a(t) = A(t)$ for all $t > 0$ gives

$$(N_s(t) + N_i(t))u'(A(t)) = (1 + \alpha)\beta A(t)(\lambda_s(t) - \lambda_i(t))N_s(t)N_i(t), \quad (18)$$

$$(\rho + \delta)\lambda_s(t) - \lambda'_s(t) = u(A(t)) + \beta A(t)^2 N_i(t)(\lambda_i(t) - \lambda_s(t)), \quad (19)$$

$$(\rho + \delta)\lambda_i(t) - \lambda'_i(t) = u(A(t)) + \alpha\beta A(t)^2 N_s(t)(\lambda_i(t) - \lambda_s(t)) - \gamma(\kappa + \lambda_i(t)), \quad (20)$$

$$\lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_s(t)N_s(t) = \lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_i(t)N_i(t) = 0. \quad (21)$$

If $\alpha = 0$, this returns the equilibrium equations (9)–(12). If $\alpha = 1$, this returns the equations (14)–(17) describing the planner's solution in the setting where there is no altruism. Intermediate values of α capture some degree of imperfect altruism in the laissez-faire setting. We note that this logic does not imply that the equilibrium is efficient under perfect altruism. The reason is that the planning problem itself changes when individuals are altruistic.

Parameters			
Parameter Description	Parameter	Value	Target
Conditional transmission prob.	β	$0.3 + \gamma$	Initial doubling time
Rate that infectiousness ends	γ	$1/7$	Duration until symptomatic
Infection fatality rate (IFR)	π	0.0062	Hall, Jones and Klenow (2020)
Value of a statistical life (VSL)	v	31,755	Hall, Jones and Klenow (2020)
Arrival rate of cure	δ	$0.67/365$	Exp. time until vaccine/cure
Discounting	ρ	$0.05/365$	Annual discount rate
Fraction initially affected	$N_i(0)$	0.0000527	Deaths before March 13, 2020
Other			
Basic reproduction number	R_0	3.1	Implied by γ and β
Expected cost of infection	κ	197	Implied by π and v
Fraction initially susceptible	$N_s(0)$	0.9999223	no social distancing before $t = 0$

Note: We calibrate the model at a daily frequency.

Table 1: Calibration to US data

5 Quantitative Analysis

5.1 Calibration

We calibrate the model to US data on the COVID-19 outbreak after March 13, 2020, the date of the proclamation declaring COVID-19 to be a national emergency in the US. Table 1 summarizes our calibration.

We normalize the time unit to be a day. We then set $\rho = 0.05/365$ to capture a 5 percent annual discount rate. In addition, we set $\delta = 0.67/365$, which implies an expected time until cure of 1.5 years. We highlight that this jointly implies that individuals heavily discount future payoffs.

Next, we set $\gamma = 1/7$, which implies that the average length of sickness is 1 week. We recognize that the disease lasts longer on average, particularly for people who develop serious symptoms. We choose this large value for γ to capture the length of the period when people are infectious and either asymptomatic, pre-symptomatic, or have mild symptoms. Lauer et al. (2020) report a median incubation period for COVID-19 of five days and that 98 percent of people who develop symptoms after an exposure do so within 11.5 days.

We calibrate β to capture an initial 30 percent daily growth rate of the stock of infected, consistent with a doubling time of approximately three days. Specifically, if $N_s(t) = A(t) = 1$, equation (8) implies $N'_i(t)/N_i(t) = \beta - \gamma$, and so $\beta = 0.3 + \gamma = 0.443$ implies an initial daily growth rate of 0.3. These values in turn imply a basic reproduction number of $R_0 = \beta/\gamma = 3.1$.¹⁰ Some recent estimates suggest an even higher value for R_0 (Sanche et al., 2020) and several authors work with a lower γ (e.g. Alvarez, Argente and Lippi, 2020). We therefore offer a robustness exercise below where we use a lower value of γ and a correspondingly higher value of R_0 , while still hitting the initial 30 percent daily growth rate.

We work with the following period utility function,

$$u(a) = \log a - a + 1. \tag{22}$$

We think of the first part as the gross returns from social activity, in particular consumption, and of the second part as the cost associated with it. This function has an interior maximum of 0 achieved at $a = a^* = 1$.

To measure the cost of disease κ , we build largely on Hall, Jones and Klenow (2020). Recall that κ is the product of the infection fatality rate (IFR) π and the value of a statistical life (VSL) v . The first term is relatively uncontroversial. Hall, Jones and Klenow (2020) work with two baseline IFRs of $\pi = 0.0081$ and 0.0044 , and so we choose the intermediate value $\pi = 0.0062$.¹¹

The VSL is a more controversial number, even though it is familiar in the economic analysis of environment regulation.¹² We follow Hall, Jones and Klenow (2020), who calculate a value of \$270,000 for each remaining year of life. They also offer a value of 14.5 years for the average remaining life expectancy of COVID-19 victims. This gives

¹⁰There is still considerable disagreement on the uncurtailed doubling time or, alternatively, R_0 of COVID-19. Johns Hopkins University’s Center for Systems Science and Engineering reports a doubling time of 2–5 days in the US in the early stages of the epidemic (Dong, Du and Gardner, 2020). Pan et al. (2020) report an effective reproduction rate slightly above 3.0 in January in Wuhan.

¹¹Robert Redfield, the director of the CDC, said during a call with reporters on June 25, 2020 that “Our best estimate right now is that for every case that was reported, there actually were 10 other infections” (<https://n.pr/3dC55bu>, retrieved June 26, 2020). With a current case count of 2,697,326 and 129,007 deaths, this would give $\pi = 0.0043$, close to the lower value in Hall, Jones and Klenow (2020) (<https://coronavirus.1point3acres.com/en>, retrieved July 1, 2020)

¹²See Greenstone and Nigam (2020) for a useful discussion of this value and the use of VSL in calculations such as ours.

$v = \$3,915,000$. Roughly speaking, this means that the typical victim of COVID-19 would pay \$3,915 to avoid a 0.1 percent probability of death. With a discount rate of $\hat{\rho}$, this is equivalent to paying a constant stream of $\hat{\rho} \times \$3,915$ per day to avoid this death risk.¹³ Compare this to current US per capita consumption of approximately \$45,000 per year (Hall, Jones and Klenow (2020)), and we reach the conclusion that people would permanently give up $\frac{\$3,915 \times 365}{\$45,000} \hat{\rho} = 31.755 \hat{\rho}$ percent of their consumption to avoid a 0.1 percent death risk.

To see how to map this into our model, ask someone with preferences as in equation (22) what fraction x of her consumption she would be willing to give up to avoid a 0.1 percent probability of death. If the answer is $31.755 \hat{\rho}$, then v solves

$$\frac{\log(1)}{\hat{\rho}} - 0.001v = \frac{\log(1 - 31.755 \hat{\rho})}{\hat{\rho}}$$

or $v = -1000 \log(1 - 31.755 \hat{\rho}) / \hat{\rho} \approx 31,755$ in our model units. The approximation is valid when $\hat{\rho}$ is small, as is the case when we use a daily discount rate. Note that this represents an “exchange rate” from model utils to US dollars of $31,755 / 3,915,000 \approx 1/123$. We use that rate when reporting on the dollar cost of the disease below.

These values are similar to the ones chosen in several other recent paper on the outbreak. For instance, Alvarez, Argente and Lippi (2020) fixed the IFR at one percent and the VSL at \$1.33 million. Furthermore, summing the age-specific VSL from Greenstone and Nigam (2020) and weighting them by the CDC estimates of age-specific death rates leads to a VSL of \$4.5 million, very similar to the number we use. We also offer a robustness exercise below where we consider reducing κ by fifty percent to reflect a lower IFR or VSL.

Finally, we calibrate the initial infection rate, $N_i(0)$, to match the 51 people who died from COVID-19 in the US on or before March 13, 2020, the date when people began to socially distance (<https://coronavirus.1point3acres.com/en>, retrieved July 1, 2020). Given a population of 328 million individuals and an IFR of $\pi = 0.0062$, this implies that $N_r(0) + N_d(0) = \frac{51}{328 \times 10^6 \pi} = 0.0000251$. We assume that prior to date 0, people were unaware of the disease and so set $A(t) = a^* = 1$ for $t < 0$.

¹³This notation allows for $\hat{\rho}$ to differ from ρ for expositional purposes, taking into account the short remaining life expectancy of the average COVID victim. It will turn out that the following calculations are (approximately) independent of the discount rate.

Following the dynamics of the SIR model, this implies $N_i(0) = 1 - (N_r(0) + N_d(0)) - \exp(-R_0(N_r(0) + N_d(0))) = 0.0000527$, leaving the remaining share 0.9999223 of the people susceptible to the disease.

We check the robustness of our results to the choices made in this section in Appendix C.1. We change preferences, lower the cost of the disease κ , increase the duration of infectivity $1/\gamma$, reduce the cure probability δ , and start with a higher stock of initially infected $N_i(0)$. For each of these exercises we contrast the paths of N_s, N_i, A, R with our baseline for both equilibrium and optimum. The main features that we discuss and highlight in the body of the paper are robust to those changes.

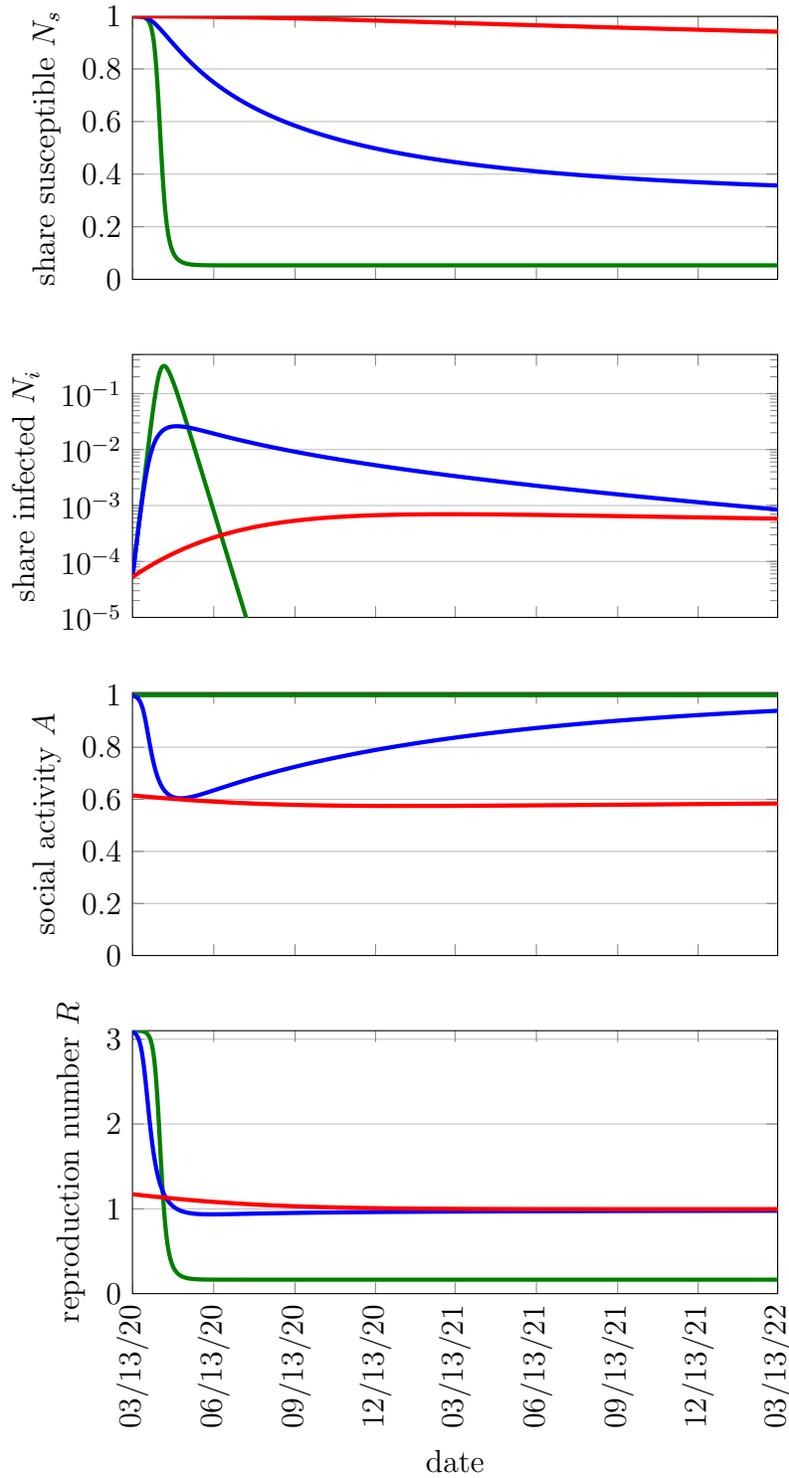
5.2 Basic SIRD Model

We next turn to our quantitative results. As a simple benchmark, we begin with the basic SIRD model and then turn to laissez-faire equilibrium and the social optimum.

The green line in Figure 4 plots the dynamics of the pandemic in the SIRD model, that is without any behavioral response, $A(t) = a^* = 1$, under the assumption that a cure is not found. The top panel shows the share of people susceptible in levels. The second panel shows the share infected on a log scale. According to this model, the pandemic would have unfolded rapidly even though the infection rate started at half a basis point in mid March. The infection rate would have peaked in the second half of April at above 31 percent. As a consequence of the rapid spread of the disease, the benefits of herd immunity would not have kicked in before almost everyone is sick. by June 1st, only 5.5 percent of the population would have remained susceptible.

This model fails to capture the experience in places that did not institute any restrictions on social activity. For instance, Sweden hit 1063 total cases, approximately 0.01 percent of the population, on March 15, 2020. One month later, this number rose to 12,443 confirmed cases, 0.12 percent of the population, despite the laissez-faire approach taken by the Swedish government.¹⁴ In contrast, our calibrated SIRD model without any behavioral response predicts that it should take about eight days for the total case rate to rise from 0.01 percent to 0.12 percent of the population. Even if the total case rate is understated by an order of magnitude, approximately exponential growth implies that it still takes only nine days for the infection rate to

¹⁴Retrieved from <https://www.worldometers.info/coronavirus/country/sweden/>.



Notes: See Table 1 for calibration. The second plot is drawn on a log scale.

Figure 4: SIRD vs. Laissez-Faire vs. Optimal Policy.

increase from 0.1 percent to 1.2 percent of the population in the SIRD model.

Likewise, SafeGraph and Google data document a sharp decline in individual social activity even in laissez-faire settings. The basic SIRD model with constant parameters misses the fact that individual behavior responds to the risk of infection.

5.3 Laissez-Faire Equilibrium

We next turn to the disease dynamics in our laissez-faire equilibrium, which we depict as the blue line in Figure 4. The difference between laissez-faire and the basic SIRD model is stark. Despite the government not intervening at all, the peak infection rate is less than one tenth of the level in the SIRD model, 2.6 percent at the beginning of May. In turn, the response of individual behavior substantially prolongs the epidemic, with the infection rate staying elevated for a much longer time. That is, equilibrium social distancing flattens the infection curve.

The third panel shows that individuals reduce their social activity by as much as 40 percent. However, the contraction in social activity does not occur immediately. Instead, it happens rapidly once the infection rate exceeds 0.1 percent. Conversely, activity starts to recover shortly after the infection rate peaks.

The fourth panel depicts the dynamics of the effective reproduction number $R(t)$. It is nearly equal to R_0 at the onset of the pandemic and then drops rapidly because of self-imposed social distancing. It falls below 1 in early May, exactly when the infection rate starts declining. Remarkably, it stays very close to 1 after this, never falling below 0.935. This is in line with the observation that N_i decays only slowly.

5.4 Optimal Policy

The red line in Figure 4 shows the optimal policy. Reflecting the external effects of social activity, one key property of the optimal policy is immediate social distancing in order to delay the spread of the infection and buy time to find a cure. While peak infection in the SIRD model occurs after 36 days and the equilibrium behavioral response delays the peak until 50 days have lapsed, the optimal policy delays it for almost a year. Because infections increase more slowly, the peak infection rate is almost forty times lower under the optimal policy than with laissez-faire, 0.07 percent.

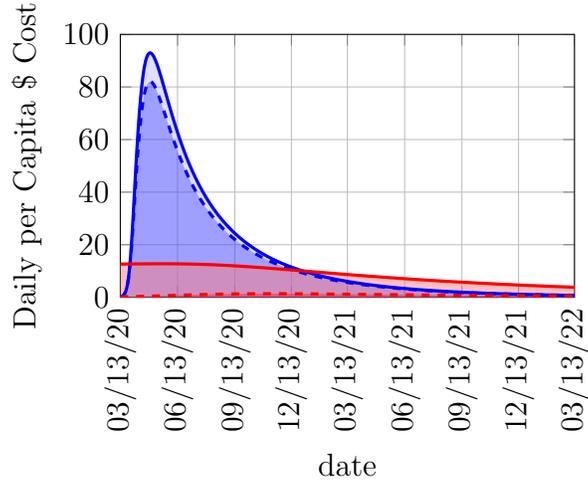
On the other hand, once the infection rate peaks, it declines extraordinarily slowly along the optimal path. One way to quantify this is through the effective reproduction number $R(t)$, which never falls below 0.996. This means the infection rate never declines by more than 0.4 percent per week. We stress that this strong desire to “flatten the curve” is true even though there is no cost of peak-loading infection rates, i.e. even though κ is constant. A healthcare capacity constraint, where the cost κ is an increasing function of N_i , would make the case for flattening the curve even stronger. We have opted to omit this from the analysis to highlight that the optimal policy aggressively flattens the curve even without such a constraint.

5.5 Cost of the Disease

We now explore the cost of the disease. In Figure 5, we plot expected flow utility in the laissez-faire equilibrium and under optimal policy. We break each of these into two components, the total health cost at any point in time, $\kappa\gamma N_i(t)$, and the total utility cost created by social distancing, $(N_s(t) + N_i(t))u(A(t))$. We then discount these back to time 0, accounting both for impatience and for the chance of a cure. Finally, we convert utils into US dollars using the exchange rate discussed in Section 5.1.

In Table 2, we integrate these curves over the infinite future to find the expected present value of the disease and its components. We do this not only in the laissez-faire equilibrium and social optimum, but also under two extreme policies: the basic SIRD model, where $A(t) = 1$ for all t by assumption; and a policy of permanently suppressing the disease, $A(t) = 1/\sqrt{R_0}$ for all t , which ensures that $R(t) = N_s(t) < 1$, so the disease always shrinks and herd immunity never emerges. The first column of the table reports the expected per capita health cost of the disease in dollars. The second column calculates the per capita dollar value of the expected reduction in social activity. Finally, the third column gives the total cost.

We first note that the composition of the cost differs sharply between the four policies. In the SIRD model, there is only a health cost by assumption. With a permanent lockdown at $A(t) = 1/\sqrt{R_0}$, the health cost is negligible since the disease never takes off. The two more interesting policies fall in between. In the laissez-faire equilibrium, about 90 percent of the cost is health, while under the optimal policy, about 90 percent is the cost of social distancing. Figure 5 shows that this



Notes: Below dashed line: total daily per capita health cost $\kappa\gamma N_i(t)$ in period t . Between dashed and solid line: total daily per capita cost from social distancing, $(N_s(t) + N_i(t))u(A(t))$. We multiply each of those with $e^{-(\rho+\delta)t}$ to discount to time 0. We convert utils to US dollars using a factor of 123 as discussed in Section 5.1. Equilibrium in blue, optimum in red.

Figure 5: Cost of Disease and its Components

decomposition does not vary dramatically over time.

Figure 5 also shows that the time path of the cost differs sharply between laissez-faire and optimum. The planner flattens the curve through persistent social distancing while the laissez-faire economy is effectively at herd immunity after 2 years. As a consequence, the expected flow costs are concentrated in the first year under laissez-faire, while a far larger fraction of the cost falls further into the future under the optimal path. At the same time, we see that the planner also incurs immediate costs from social distancing in excess of those that would occur in equilibrium.

Next, we note that the disease is very costly. The expected discounted flow cost per capita even under the optimal policy starts at \$13 per person per day and is still at \$4 per day two years later, with much of the decline reflecting the 74 percent possibility of a cure within that horizon. Under laissez-faire, the peak cost is \$93 per day by early May, but then it declines rapidly. By the start of 2021, the expected daily flow costs under laissez-faire fall below those of the optimal path.

This observation is evident in aggregate costs as well. Again using \$45,000 as annual per capita consumption, the expected present value of the cost of the disease amounts to over 28 percent of annual per capita consumption under laissez-faire and

	Health	Activity	Total
SIRD	\$21,223	0	\$21,223
Laissez-Faire	\$11,374	\$1,277	\$12,652
Optimum	\$860	\$7,227	\$8,087
$A(t) = \frac{1}{\sqrt{R_0}}$	\$78	\$8,334	\$8,412

Notes: First column: present value of per capita pure health cost of disease. Second column: present value of per capita cost of reduction in social activity. Third column: total cost. First row is the SIRD model, second row is the equilibrium, third row is the optimum, and the last row is a permanent lock down from the outset that never allows the disease to take off, $A = \frac{1}{\sqrt{R_0}}$. We convert utils to US dollars using a factor of 123, as discussed in Section 5.1.

Table 2: Per Capita Welfare Cost of Disease

18 percent under the optimal policy. While these numbers are large, they are much smaller than the naïve calculation without a behavioral response would suggest, a health cost equal to 47 percent of annual per capita consumption. Moreover, these numbers seem reasonable in light of the events we have seen in the Spring of 2020. As we write this paper, almost 130,000 people have died, while the US unemployment rate has declined only slightly from record levels, with second quarter GDP forecast to contract by almost fifty percent.¹⁵

The benefits from pursuing the optimal policy are large. Comparing the aggregate costs in Table 2, the planner is able to cut expected per capita costs by 36 percent compared to laissez-faire. This reflects that the negative externalities of social activity in the COVID-19 outbreak were large, making the case for mandatory social distancing was strong. On the other hand, an extreme policy of cutting social activity to $1/\sqrt{R_0}$ until a cure is found would reduce the mortality cost by another order of magnitude but would cause too much economic hardship, pushing total costs 4 percent above the optimal level. The fact that the gains from the optimal policy are small compared to what is achieved by a strict lockdown reflects the fact that the optimal time path for social activity is very flat (Figure 4).

Finally, we can express the health cost in terms of fatalities. Health costs amount to 0.29 percent of the value of a statistical life under laissez-faire and 0.022 percent under the optimal policy. Given a US population of 328,000,000 people, this makes the health cost of laissez-faire equilibrium equivalent to around 953,000 immediate

¹⁵Atlanta Fed NowCast, <https://tinyurl.com/yc46b392>, retrieved June 26, 2020.

deaths, vastly more than the 72,000 immediate deaths associated with the optimal policy. Put differently, the expected economic cost of \$5,950 incurred by each US citizen due to mandatory social distancing rules would have saved over 881,000 lives in expectation. Of course, the US has already experienced more than 72,000 immediate deaths by mid-2020, which suggests that the implemented policies were not optimal.

An important caveat is that these calculations correspond to the *expected* present value cost, where uncertainty comes from the timing of the cure. If a cure comes very late, the *realized* present value will be larger under the optimal policy. This reflects the sharp difference in the time path of costs, with optimal policy delaying the disease in the hope of a cure. If there is no cure, then fatality rates are similar under the optimal policy and in laissez-faire, while the optimal policy imposes far greater costs of social distancing.

In the robustness Appendix C.1 we show how the level and composition of the cost of the pandemic vary with substantial changes to our calibration. We find that the laissez-faire results are quite robust, with health costs always accounting for more than 80 percent of the total. The gains from, and cost compositions of, optimal policy are more sensitive to the calibration. When the gains from social distancing are smaller because a cure is less likely or health costs are lower, we find that optimal policy allows an initial wave of infections. The same thing happens when the basic reproduction number is higher. In this case, effective social distancing requires a bigger reduction in activity, which the planner finds too costly. This means that in those cases, the laissez-faire policy is less inefficient and that the health costs of the two policies are more similar.

6 Model versus Data

This section compares our model's predictions with evidence on social activity, on the effective reproduction number, and on mortality in order to evaluate our model's quantitative predictions and interpret Spring 2020 through the lens of the model.

6.1 Social Activity

We first re-visit the quantitative evidence from Google and SafeGraph from Section 3 and contrast it with our model’s implications for social activity. We note that we did not target the response of social activity to the COVID-19 outbreak in our calibration. We do not claim that the hastily-implemented lockdowns and mobility restrictions were close to the social optimum, but we believe that the response in individual behavior witnessed prior to implementation of any policy measures should be picked up by the laissez-faire equilibrium.

Figures 1, 2, and 3 show a 40 to 50 percent decline in activity across the different metrics and data sets we study. In the US, much of this decline occurred before stay-at-home orders were issued and the decline was remarkably uniform across a variety of locations. Since the empirical metrics we have studied all have an inherent cardinality, their decline is quantitatively meaningful. Moreover, we can also compare these metrics to the decline in $A(t)$ in the laissez-faire equilibrium. As can be seen in Figure 4, individual activity declines by up to 40 percent in equilibrium. Thus, the model captures the magnitude of the decline in activity quite well.

On the other hand, while the speed of the decline in activity in the model is fast, the speed of decline in the data was even faster. In the model, $A(t)$ falls by 1.6 percent during the week of March 13–20 and takes three more weeks to fall by 31 percent, while in the data we observe a decline of that magnitude during the first week, before stay-at-home orders were implemented. This may reflect the possibility that individuals overestimated the risk of infection during this first week, a reasonable possibility given the scarcity of data and abundance of rumors. Alternatively, it may reflect altruistic behavior, which made the initial decline more in line with the red line in Figure 4, i.e. an immediate 39 percent drop in activity as soon as people became aware of the disease.

We could have directly targeted this type of data in our calibration, for example to select a value for the cost of the disease κ . That is, a natural approach would be to let the individual behavioral response “reveal” the perceived cost of infection instead of relying on direct measures of the infection mortality rate, π , and the value of a statistical life, v . Instead, we have opted with the direct measures which allows us to use the data on individual pre-lockdown behavior as model validation.

6.2 Effective Reproduction Number: Roll-over after Peak

A robust feature of our analysis is that the effective reproduction number $R(t)$ remains close to 1 after crossing this threshold. This implies infections decline very slowly both under laissez-faire and under the optimal policy. We now contrast this feature with two direct empirical counterparts. Fernández-Villaverde and Jones (2020) develop a method to invert the SIR model and estimate $R(t)$ from time series on deaths alone. They do so for 88 countries, states, and cities (not mutually exclusive); see https://web.stanford.edu/~chadj/JFV_Jones_DetailedResults.xlsx.¹⁶ We complement this with an alternative measure of $R(t)$ provided by the website rt.live. This uses time series on testing volumes and positive tests, rather than deaths. It offers separate measures for all 50 US states and the District of Columbia.¹⁷

To focus on the behavior of $R(t)$ around peak infections and after we proceed as follows. For each location in the respective data set, we identify the date when the estimated $R(t)$ falls below 1 for the first time following the outbreak. We then track the evolution of $R(t)$ around that date (from 25 days prior to 60 days past) and plot it in Figure 6 in various different ways.

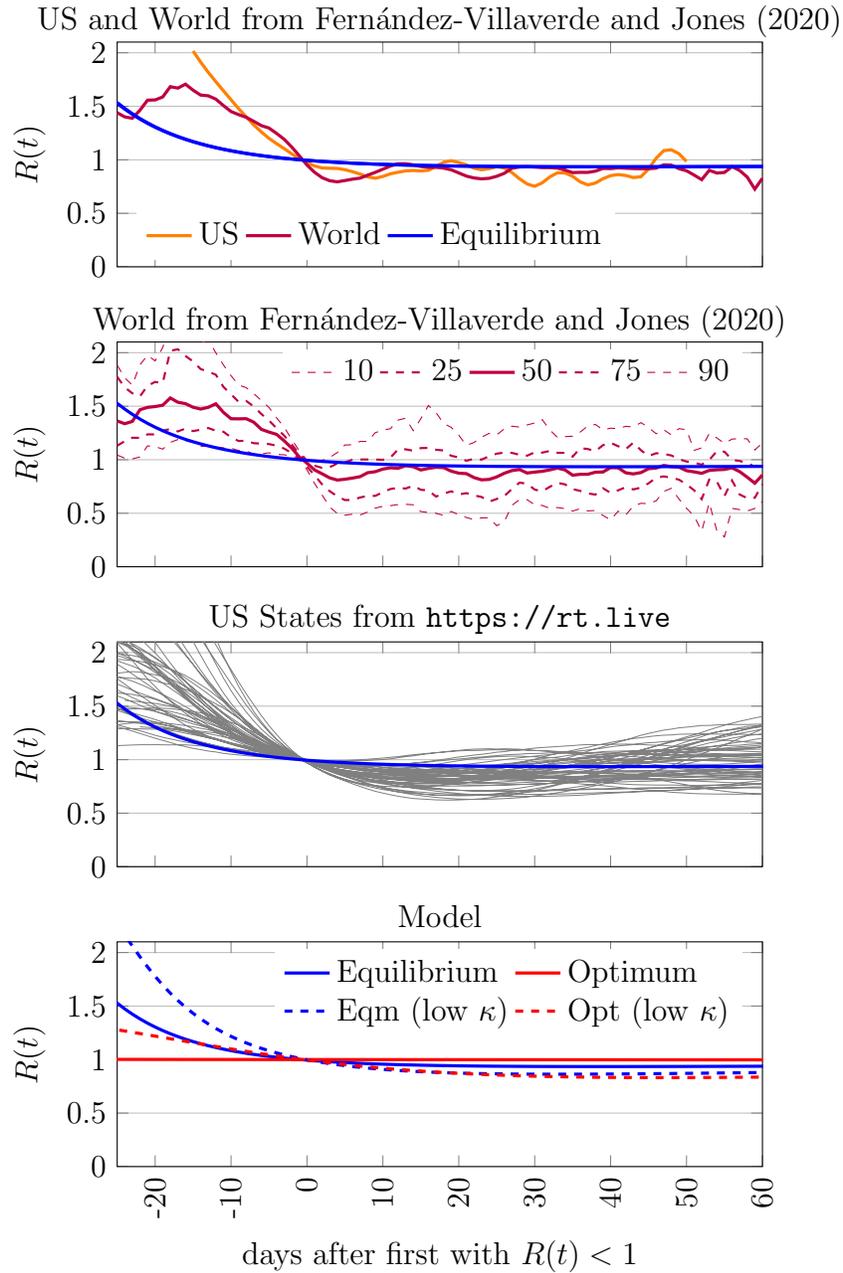
The top panel plots the evolution of $R(t)$ for both an unweighted average across all locations studied by Fernández-Villaverde and Jones (2020) and for the US only. We benchmark this with the equilibrium version of our model. The evolution of $R(t)$ after aligns remarkably well with our laissez-faire model. It appears like both the US and the world average have been rolling over infections much like our model predicts.

The second panel plots the same across all 87 locations studied by Fernández-Villaverde and Jones (2020). We plot, for each day around peak infections, the following percentiles of $R(t)$ across locations: 10, 25, 50, 75, 90. Again, the median aligns closely with the model. There is substantial dispersion across locations but we note that the data contains locations such as Armenia, Estonia, Moldova, Andorra, along with many US states and cities; those have small populations and imperfect data so some noise is to be expected.

The third panel proceeds identically but plots the data provided by rt.live for

¹⁶We drop the time series for Hubei, China which ends on Feb 19th and only covers 24 days.

¹⁷Data retrieved from download link on <https://rt.live>, <https://d14w1fuexuxgcm.cloudfront.net/COVID/rt.csv>, on 6/24/2020.



Notes: For location specific results we use results from Fernández-Villaverde and Jones (2020) and rt.live who estimate location-specific time series for $R(t)$. For the plot we identify, for each of the location in the data-set, the first day $R(t)$ falls below 1 following the COVID-19 outbreak and plot $R(t)$ 25 days prior and 60 days past that date. Model counterpart corresponds to baseline calibration, except in the bottom panel, where the dotted lines reduce the cost of disease κ by half, as in Figure 12.

Figure 6: Effective reproduction number $R(t)$ around peak infections.

all 50 US states plus the District of Columbia. All 51 locations display the same quantitative pattern. After peak infections, $R(t)$ stays in close vicinity of 1.

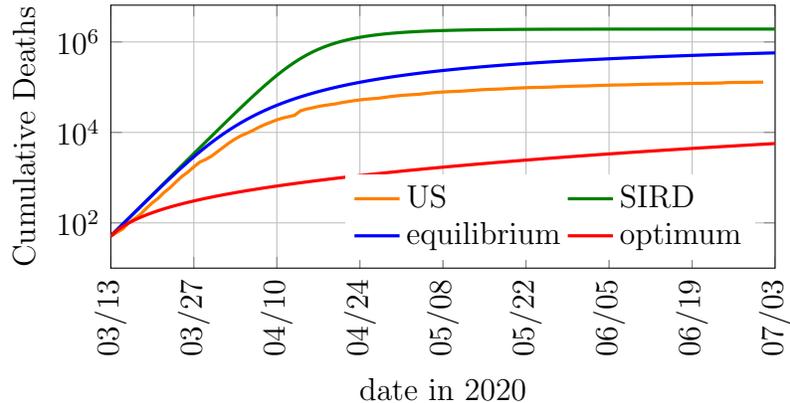
In the bottom panel, we plot the evolution of $R(t)$ around peak not only in the laissez-faire version of our setup but also under the optimal policy. What is clear from the figure is that, under the optimal policy just as under laissez-faire, $R(t)$ stays very closely below 1 after peak infections. Again, this feature of the optimal policy is very robust and arises for the same reason. The state variables are moving slowly, which makes optimal behavior and policy sticky and implies the roll-over feature.

However, optimal policy prior to peak infections is much “flatter” under the baseline calibration and, as a consequence, does not align with the observed time series. As mentioned, it is hard to judge to what extent the policy measures taken in the Spring of 2020 correspond to the optimal policy. However, we highlight that this mismatch is not generic to the model. Indeed, if the expected cost of disease is 50 percent lower (a case we study in Appendix C.1 and depict in Figure 12), then the optimal path for $R(t)$ prior to peak infections aligns more closely with the data.

Taken together, the figure shows that data on the effective reproduction number aligns with our setup. In particular, the data are consistent with the roll-over feature that we have highlighted in the model. $R(t)$ remains closely below 1 for as long as we currently have data in the US and many other locations. The dispersion is surprisingly small given the wide range of locations studied and the inherent difficulty in measuring $R(t)$, which creates measurement error. We show in Appendix C.1 that the finding that $R(t)$ never falls far below 1 after the infection rate peaks holds with substantially different calibrations of the model; see the bottom panel in Figures 11–15. In Section 7, we provide some intuition for this result.

6.3 Fatalities

We conclude this section by contrasting the model with data on US COVID-19 fatalities. In Figure 7, we plot the cumulative fatality rate in our model, which is just $\pi(1 - N_s(t) - N_i(t))$, multiplied by the US population of 328 million people. We plot this for the basic SIRD model along with the laissez-faire and optimal policy paths of our model. All four plots start with total COVID-19 fatalities for the US on March 13, 2020. According to the SIRD model, 1,925,000 Americans would have



Notes: Cumulative fatalities (without a cure) in the model and data. US data comes from <https://coronavirus.1point3acres.com/en>, retrieved on July 1, 2020. We calibrate the model to match data on cumulative deaths on March 13.

Figure 7: Cumulative Fatalities

died by July 1st from the disease without any social distancing. Equilibrium behavior reduced that by 70 percent, to 564,000. Actual deaths were yet another 77 percent lower, 129,000, showing the importance of some combination of altruism and mandated social distancing. On the other hand, the optimal policy would have crushed this outcome, with only 5,489 dead.¹⁸

7 Key Qualitative Features

In what follows we summarize a few key properties that are either common between equilibrium and optimum, or help to distinguish them. This discussion enables us to better frame the observed empirical patterns in the context of our model. We first highlight two features which distinguish optimum from the laissez-faire, and then move to characteristics which are more similar.

Immediate Social Distancing The optimal policy delays the onset of the disease by imposing immediate social distancing, as depicted in the third panel of Figure 4. In particular, the bottom panel shows that the optimal effective reproduction number

¹⁸It is worth noting that this mortality rate, 0.002 percent, is lower than Germany's (0.011 percent) but higher than South Korea's or New Zealand's (both 0.0005 percent).

is initially far below its uncurtailed counterpart $R_0 = 3.1$. Even if a full outbreak is eventually inevitable, the social gains from immediate social distancing are of first order because of discounting and because of the hope for a cure. The cost, however, are of second order, since $u'(a^*) = 0$.¹⁹

On the other hand, there is no noticeable reduction in social activity under laissez-faire at the initial outbreak of the disease. As a result, the laissez-faire effective reproduction number is indistinguishable from R_0 in the early days of the disease. This reflects the fact that there is little private incentive to lower social activity when the risk of individual infection is negligible, and so individuals do not restrict activity until infections are rampant.

Persistent Social Distancing Under laissez-faire, social activity is almost back to its pre-pandemic level after 2 years. This is not the case under the optimal solution, which curtails activity for decades or until a cure is found. That is the flip-side of delay. The optimal solution delays the onset of the disease in the hope of finding a cure. If no cure is found, herd immunity only builds very slowly and so restrictions on social activity must persist far longer than under laissez-faire.

Long Run We now highlight three observations about the long run, common to both equilibrium and the social optimum.²⁰ Of course, those are only relevant in the case where no cure is found for a long time.

First, the disease vanishes in the long run. This reflects three assumptions. First, the population is finite and there are no newly-born individuals added to the population. Second, individuals attain permanent immunity after infection. Without either of these assumptions, an “interior steady state” with a strictly positive stock of infections would naturally emerge. Third, the probability that an individual stays sick for t periods converges to 0 as t grows large, a natural assumption.

The next observation is that social distancing vanishes asymptotically with social

¹⁹In contrast, Alvarez, Argente and Lippi (2020) assume the population without disease is in a corner and so reductions in social activity, lockdowns in their words, have a first order cost.

²⁰Since we only plot the time series for 2 years these are hard to see for the planner. These are, however, numerically accurate observations. We solve the model until the infection rate falls to $N_i(T) = 10^{-9}$, which takes 569,000 days, or sometime in the year 3578. The probability of not finding a cure by then is vanishingly small.

activity returning to its pre-pandemic level a^* . This is not obvious ex-ante because it is feasible to suppress social activity permanently in order to further improve health outcomes, for example by setting $A(t) = 1/\sqrt{R_0}$. To gain some intuition for this result, suppose the planner pursued a path of permanently suppressing social activity. Eventually, allowing a bit more social activity becomes arbitrarily cheap because the number of infected people is so small that any meaningful outbreak of the disease is pushed far off into the future. In contrast, there is a first order gain from the increase in social activity because there are always susceptible individuals suffering from the suppression of activity.

The key assumption here is that the disease cannot be eradicated through social distancing, which seems a plausible assumption. For example, even if the disease is temporarily pushed out of the human population, there is evidence that it can survive in animals and then reenter the human population through zoonotic transmission.

An immediate consequence of vanishing social distancing is that the population eventually reaches herd immunity, with $N_s(t)$ falling below $1/R_0$. This is because when social distancing vanishes, $A(t) \rightarrow 1$, infections shrink if and only if $N_s(t) < 1/R_0$; see equation (8). Put differently, vanishing social distancing and a disappearing disease are only mutually consistent if there is herd immunity in the long-run.

The last two observations are closely intertwined. The disease does not vanish due to permanent social distancing, but due to herd immunity. We stress that there are feasible paths where non-negligible levels of social distancing remain until there is a cure, and so herd immunity is never achieved. Such paths are not optimal. Of course, the optimal path towards herd immunity is extremely slow and, given that a cure eventually puts an end to the disease, one that is highly unlikely to fully unfold.

Mild Social Distancing: the Effective Reproduction Number At the onset of the pandemic, the effective reproduction number $R(t)$ differs significantly between the laissez-faire equilibrium and social optimum. Specifically, under the optimal path it drops immediately due to social distancing, while the decline is more continuous under laissez-faire. But under both policies, $R(t)$ eventually crosses 1 and after this point, the two paths are remarkably similar, with $R(t)$ never falling far below 1, effectively rolling over the stock of infections.

To understand why $R(t)$ never falls much below one, it helps to consider the evolution of the state variables $N_s(t)$ and $N_i(t)$. Consider first the point when infections reach their peak, that is when $R(t) = 1$. At this stage, the state $N_i(t)$ does not move by definition. One sick person gets replaced by another. As a consequence, the only reason this is not actually a steady state is that the pool of susceptibles drains, $N'_s(t) < 0$. But since under our calibration $N_i(t)$ is never very high, $N_s(t)$ moves slowly. The slow moving state variables imply optimal behavior is also slow moving, and so $R(t)$ must stay very close to 1 even after the peak in infections.

We add two observations to this. First, $N_i(t)$ is higher and hence $N_s(t)$ moves faster under laissez-faire than under the optimal policy. This is why $R(t)$ falls further under laissez-faire compared to the optimum. Second, suppose the peak level of infections was much higher, as would be the case if the cost of the disease were low. In this case, $N_s(t)$ would be fast moving and we would see a bigger decline in $R(t)$ after the peak of infections. We verify this in Figures 12–14 in Appendix C.1. In each figure, we change some parameter to increase the peak infection rate. $R(t)$ declines by more, but the decline is modest in each case.

Another perspective is that there is an important stabilizing force pushing $R(t)$ towards 1. When $R(t)$ is larger than 1, the number of infections grows, discouraging social activity. But a lower level of social activity mechanically reduces $R(t)$. The reverse happens when $R(t)$ is smaller than one. The only reason $R(t)$ does not converge to 1 is because of the decline in the number of susceptible people. This also works to make individuals more cautious, reducing $A(t)$ and hence $R(t)$. This is why $R(t)$ slightly overshoots 1.

Finally, we offer another way to think about the finding that social distancing is never too extreme. In both the laissez-faire equilibrium and optimal policy, the level of social activity $A(t)$ never falls below $1/\sqrt{R_0}$, the critical value below which social distancing would wipe out the disease for any $N_s(t) \leq 1$. This result is robust to a wide range of parameters. The only exception is if the initial infection rate is large. Even in this case, we have $A(t) < 1/\sqrt{R_0}$ initially, but social activity later rises above this threshold and never falls below it again.

8 Discussion of Key Assumptions

Information Structure We have assumed that individuals do not know they are sick until they recover or die. Most of the papers that we cite in Section 2 assume that individuals always know their health status. As a consequence, individuals ramp up their social activity (or consumption and labor supply) at the moment that they become sick because they have nothing left to lose. We do not believe that this is realistic, both because many individuals are pre-symptomatic or asymptomatic after they are infected, and because many people try not to get others sick.

In any case, if individuals know whether they are sick, eliminating the disease becomes both feasible and cheap with appropriate policy instruments. A planner that can directly gear policy towards the infected can isolate them and curb the spread of the disease at little cost. In practice, however, this only becomes relevant once widespread test-and-trace (and possibly quarantine) is available and functional.

We have also assumed that the recovered know that they have gotten infected. This is inconsistent with the observation that many infected people never develop any symptoms. It is also inconsistent with the observation that the recovered were subject to the same restrictions on social activity as other people. Assuming the opposite has two opposing effects. On the one hand, social distancing is an even costlier tool of disease prevention because it applies to a larger population, especially in the later stages of the pandemic when many people have recovered. On the other hand, since social activity is still suppressed for recovered people, acquiring immunity through infection is somewhat less attractive. In Appendix C.2, we work through the math of the model and characterize both the laissez-faire policy and the social optimum. We find that changing this assumption does not have a big quantitative impact on either the laissez-faire policy or the optimal policy. Under laissez-faire, social activity already reverts back to normal fairly quickly after infections peak. It reverts back slightly faster when individuals do not realize they have recovered. Under the optimal policy, few people get sick and so the recovered population is of negligible size. The planner suppresses activity slightly more.

Available Policy Instruments The policy maker in our setting only has access to a very blunt policy tool: widespread social distancing. Perhaps most importantly,

there is no test, trace, and quarantine alternative to complement social distancing. This is a natural benchmark given the information structure we just discussed and given the realities of the pandemic thus far in the US and much of Europe.

The availability of a test, trace, and quarantine policy could alter our conclusions because eradicating the disease may become optimal. The key observation is that when the share of infected $N_i(t)$ becomes small, widespread social distancing is a blunt tool because it primarily affects the large pool of susceptible individuals, $N_s(t)$. On the other hand, test, trace, and quarantine scales with the size of the infected population and so the cost vanishes as the disease vanishes. The important takeaway is that widespread social distancing is not a tool that should be used with the aim of eradicating the disease.

Healthcare Capacity Constraint We have assumed that the probability of death π and hence the cost of the disease κ are independent of the number of infected people $N_i(t)$. This ignores the possibility that a sudden surge in infections may overwhelm the healthcare system. This is not because we want to deny the relevance of healthcare capacity constraints. They were clearly an important consideration during the early weeks of the COVID-19 outbreak; see <https://twitter.com/drewaharris/status/1233267475036372992>. Instead, our point is that economic logic would push us to flatten the curve even if there were no healthcare capacity constraints.

Heterogeneity The fatality risk of COVID-19 is far higher for older individuals and those who have certain pre-existing conditions. Our setup abstracts from such heterogeneity. We believe that this picks up the realities of policy during the pandemic, which was applied across the board in a non-discriminatory fashion. Of course, for this reason our laissez-faire setting might look particularly bad since it does not take into account that those at a high risk would naturally reduce activity more and those with little risk less so. Nonetheless, we do not think that our key takeaways with respect to voluntary distancing, optimal policy, and the evolution of $R(t)$ after peak infections would be substantially altered in a setting with heterogenous individuals.

Another source of heterogeneity lies in the propensity to spread the disease. In particular, epidemiologists have commented on the importance of super-spreaders for COVID-19. For example, Endo et al. (2020) estimates that about ten percent of

infected people account for 80 percent of new infections. If doing so were feasible, it would be optimal to focus social distancing on those people with a high propensity to be super-spreaders. Since it is unclear whether such people can be identified before the super-spreading events, the alternative is to focus social distancing policies on the activities where super-spreading is most likely, such as prisons, elder care facilities, and bars.

Deterministic Arrival of a Cure We have assumed that a cure for the disease arrives stochastically at rate δ . Alternatively, we can solve the model when the cure arrives deterministically at date T . This has little qualitative impact on our results in the initial phase of the disease. In particular, A , N_s , N_i , and R all behave very similarly to our benchmark, both in equilibrium and optimum. As T gets closer, however, social activity rises since the cost of infection falls to zero at the terminal date. This leads to a large increase in infections both in equilibrium and optimum, an implication that strikes us as problematic given the uncertainty inherent in medical progress.

Vaccine versus Cure Finally, we have experimented with both a stochastic and a deterministic arrival rate of a vaccine, rather than a cure. A vaccine effectively shifts people from susceptible to recovered but does not protect infected people. This alternative has no interesting quantitative impact on our results.

9 Conclusion

This paper uses a standard dynamic economic model to integrate privately optimal behavior and policy analysis into an epidemiological model of COVID-19. The model speaks to data on social activity and does a good job of describing observed behavioral patterns after the US woke up to the dangers posed by the disease in mid-March 2020.

Our quantitative exercises reveal several robust patterns. Even in *laissez-faire*, individuals significantly lower social activity in order to reduce the risk of getting infected. This behavior does not internalize the risk of getting other people sick, however. An optimal policy that internalizes this cost immediately curtails social activity to buy time. Still, social distancing is mild in the sense that it does not try to eradicate the disease. Instead, if there is no effective treatment, optimal social

distancing only vanishes in the very long run, once the population achieves herd immunity. Along the way, it is optimal to roll over a small stock of infections, which means that the effective reproduction number stays very close to 1.

Our framework is general and tractable. Equilibrium behavior and optimal policy are encoded in a set of two differential equations that can jointly be solved with the epidemiological block of the model. We therefore view the tools offered here as a natural building block to explore the role of alternative policies. The model can also be extended to analyze issues like heterogeneity and to explore the role of geography in determining disease transmission.

References

- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston**, “Optimal targeted lockdowns in a multi-group SIR model,” 2020.
- Alvarez, Fernando, David Argente, and Francesco Lippi**, “A Simple Planning Problem for COVID-19 Lockdown,” March 2020.
- Atkeson, Andrew**, “What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios,” Technical Report, National Bureau of Economic Research 2020.
- Barro, Robert J, José F Ursua, and Joanna Weng**, “The Coronavirus and the Great Influenza Epidemic-Lessons from the coronavirus potential effects on mortality and economic activity,” 2020.
- Bethune, Zachary and Anton Korinek**, “COVID-19 Infection Externalities: Pursuing Herd Immunity or Containment?,” Technical Report, Working Paper 2020.
- Bognanni, Mark, Doug Hanley, Daniel Kolliner, and Kurt Mitman**, “Economic Activity and COVID-19 Transmission: Evidence from an Estimated Economic-Epidemiological Model,” July 2020.
- Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt**, “An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies,” 2020.

- Budish, Eric**, “ $R < 1$ as an Economic Constraint: Can We “Expand the Frontier” in the Fight Against Covid-19?,” April 2020.
- Dewatripont, Mathias, Michel Goldman, Eric Muraille, and Jean-Philippe Platteau**, “Rapid identification of workers immune to COVID-19 and virus-free: A priority to restart the economy,” Technical Report, Discussion paper, Universit Libre de Bruxelles 2020.
- Diamond, Peter A.**, “Aggregate demand management in search equilibrium,” *Journal of Political Economy*, 1982, 90 (5), 881–894.
- **and Eric Maskin**, “An Equilibrium Analysis of Search and Breach of Contract, I: Steady States,” *The Bell Journal of Economics*, 1979, 10 (1), 282–316.
- Dong, Ensheng, Hongru Du, and Lauren Gardner**, “An interactive web-based dashboard to track COVID-19 in real time,” *The Lancet infectious diseases*, 2020.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt**, “The Macroeconomics of Epidemics,” April 2020.
- Endo, Akira, Sam Abbott, Adam J Kucharski, Sebastian Funk et al.**, “Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China,” 2020.
- Fenichel, Eli P**, “Economic considerations for social distancing and behavioral based policies during an epidemic,” *Journal of Health Economics*, 2013, 32 (2), 440–451.
- Fernández-Villaverde, Jesús and Chad I. Jones**, “Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities,” May 2020.
- Garibaldi, Pietro, Espen R. Moen, and Christopher A. Pissarides**, “Modelling contacts and transitions in the SIR epidemics model,” Technical Report, CEPR Covid Economics Working Paper, Issue 5 2020.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and Jose Victor Rios-Rull**, “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic,” Technical Report, Working Paper 2020.

- Goolsbee, Austan and Chad Syverson**, “Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020,” 2020.
- Greenstone, Michael and Vishan Nigam**, “Does Social Distancing Matter,” March 2020.
- Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michèle Tertilt**, “An equilibrium model of the African HIV/AIDS epidemic,” *Econometrica*, 2019, 87 (4), 1081–1113.
- Hall, Robert E., Charles I. Jones, and Peter J. Klenow**, “Trading off Consumption and COVID-19 Deaths,” 2020.
- Jones, Callum, Thomas Philippon, and Venky Venkateswaran**, “Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home,” 2020.
- Kaplan, Greg, Benjamin Moll, and Gianluca Violante**, “Pandemics According to HANK,” Technical Report, Working Paper 2020.
- Keppo, Jussi, Marianna Kudlyak, Elena Quercioli, Lones Smith, and Andrea Wilson**, “For Whom the Bell Tolls: Avoidance Behavior at Breakout in COVID19,” Technical Report, Working Paper 2020.
- Kermack, William Ogilvy and Anderson G McKendrick**, “A contribution to the mathematical theory of epidemics,” *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 1927, 115 (772), 700–721.
- Kremer, Michael**, “Integrating behavioral choice into epidemiological models of AIDS,” *The Quarterly Journal of Economics*, 1996, 111 (2), 549–573.
- **and Charles Morcom**, “The effect of changing sexual activity on HIV prevalence,” *Mathematical biosciences*, 1998, 151 (1), 99–122.
- Krueger, Dirk, Harald Uhlig, and Taojun Xie**, “Macroeconomic Dynamics and Reallocation in an Epidemic,” Technical Report, CEPR Covid Economics Working Paper, Issue 5 2020.

- Lauer, Stephen A, Kyra H Grantz, Qifang Bi, Forrest K Jones, Qulu Zheng, Hannah R Meredith, Andrew S Azman, Nicholas G Reich, and Justin Lessler**, “The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application,” *Annals of internal medicine*, 2020.
- Pan, An, Li Liu, Chaolong Wang, Huan Guo, Xingjie Hao, Qi Wang, Jiao Huang, Na He, Hongjie Yu, Xihong Lin, Sheng Wei, and Tangchun Wu**, “Association of Public Health Interventions With the Epidemiology of the COVID-19 Outbreak in Wuhan, China,” *JAMA*, 05 2020, *323* (19), 1915–1923.
- Philipson, Tomas J.**, “Economic epidemiology and infectious diseases,” *Handbook of Health Economics*, 2000, *1* (B), 1761–1799.
- **and Richard A. Posner**, *Private choices and public health: The AIDS epidemic in an economic perspective*, Harvard University Press, 1993.
- **and –**, “The microeconomics of the AIDS epidemic in Africa,” *Population and Development Review*, 1995, pp. 835–848.
- Piguillem, Facundo and Liyan Shi**, “The optimal covid-19 quarantine and testing policies,” Technical Report, Einaudi Institute for Economics and Finance (EIEF) 2020.
- Rowthorn, Bob RE and Flavio Toxvaerd**, “The optimal control of infectious diseases via prevention and treatment,” Technical Report 2020.
- Sanche, Steven, Yen Ting Lin, Chonggang Xu, Ethan Romero-Severson, Nick Hengartner, and Ruian Ke**, “Early Release-High Contagiousness and Rapid Spread of Severe Acute Respiratory Syndrome Coronavirus 2,” 2020.
- Toxvaerd, Flavio**, “Rational disinhibition and externalities in prevention,” *International Economic Review*, 2019, *60* (4), 1737–1755.
- , “Equilibrium Social Distancing,” Technical Report 2020.

Online Appendix

A Additional SafeGraph Results

A.1 Time at Home

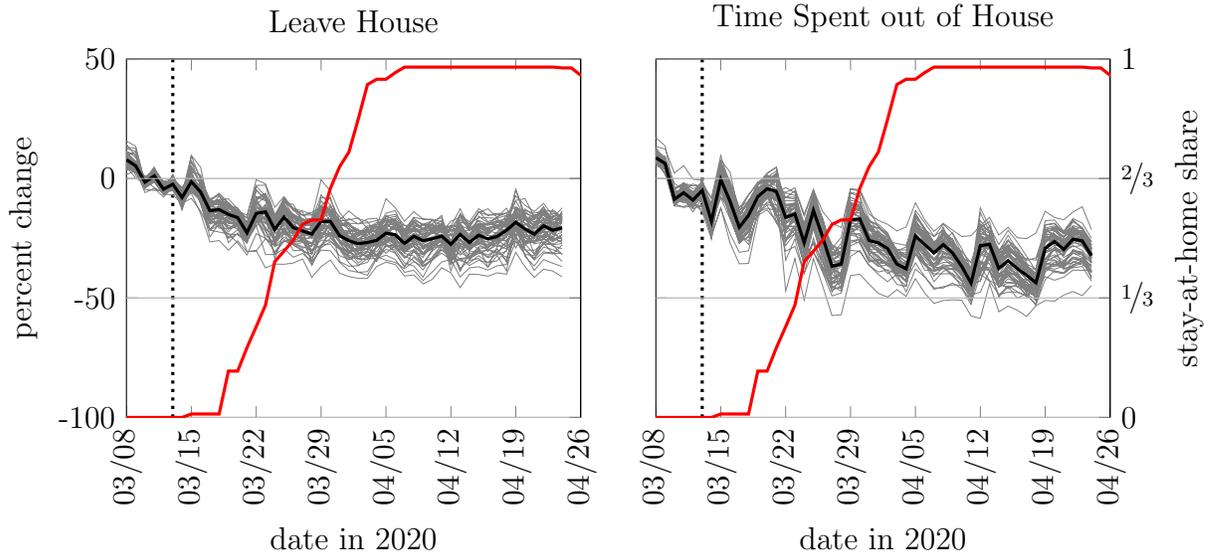
SafeGraph has a second data set named “Social Distancing Metrics.” It assigns smartphones a home address using their night-time location and then considers individual foot traffic. The dataset goes back to January 1, 2020 and currently runs until April 24. Again, it is large. On March 1, the dataset contains information from over 20 million devices across 220,000 census block groups with at least 5 devices. Among other things, it reports, for each census block group, the median number of minutes a device dwells at its home location (variable `median_home_dwell_time`). In addition, it also measures the number of devices that spend the entire day at the home location (variable `completely_home_device_count`).

We proceed in the same fashion as with our main measure. We subtract the median minutes spent at home from $24 \times 60 = 1440$ and take a daily state-wide average across census block groups using the device count as weights.²¹ We similarly construct the state-wide fraction of all devices that leave the house at least once during any day. Both are expressed relative to a benchmark of the first week of March. Figure 8 plots the evolution of these two metrics the same way as Figure 3 does for our main variable.

A.2 Closures

We work with the following closure criterion. We diagnose a POI as closed on the first of three consecutive days where daily traffic never exceeds 20 percent of the average daily traffic during a baseline period (first week of March). We only use POIs with 10 or more average daily visitors during baseline which are covered during all days the dataset covers and restrict attention to NAICS codes 44, 45, 71, 72 (retail trade; arts, entertainment, and recreation; accommodation and food services). Proceeding

²¹There are very few observations with median dwell time larger than 1440. We omit those from the analysis.



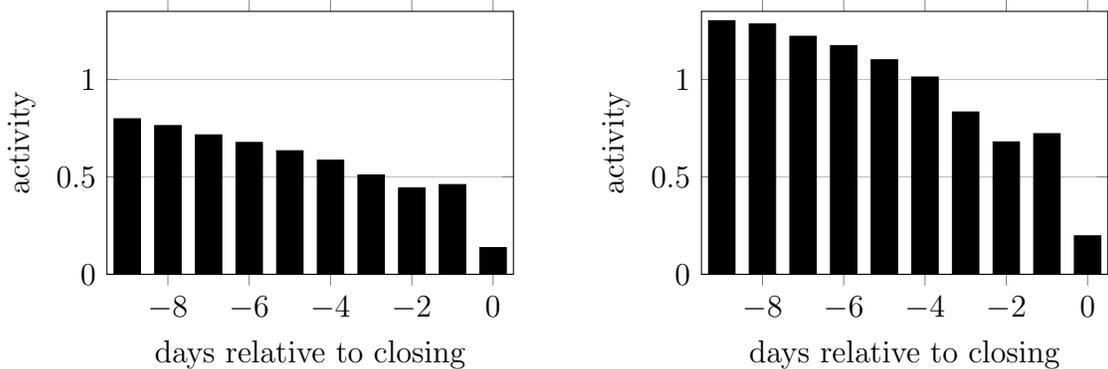
Notes: Social activity based on SafeGraph’s “Social Distancing Metrics” COVID-19 Response Dataset. Left Panel Metric: Fraction of devices that leave assigned “home” at least once during any day. Right Panel Metric: One minus fraction of time spent at home location of median device. Metrics computed at census block level, then averaged to the state level (with device-count as weight). Each thin line corresponds to one of the 51 US states and the District of Columbia. The thick black line marks the median state in terms of the decline relative to baseline at any given day. The solid red line is the percent of the population subject to stay-at-home or shelter-in-place orders, indicated on the right-hand axis. The population subject to these orders is based on authors’ own calculations using <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html>. The vertical dotted line indicates March 13, 2020, the date of the declaration of national emergency.

Figure 8: Declining Activity, Early and Everywhere IV

this way, we find that 43 percent of the 655,185 POIs had closed by the end of our sample (May 16).

In Figure 9 (left panel) we plot the average daily traffic (relative to same day-of-week traffic during baseline) in the 10 days prior to store closure. The main takeaway is that, prior to store closure revealed in this fashion, foot traffic to POIs declined substantially. This suggests that closures are not the driving force of the decline in activity that preceded the official lockdown measures.

One concern is that the pre-closure decline may reflect the closure of other businesses or offices. To address this, we separately restrict attention to POIs that receive, in the baseline week, over 50 percent of their traffic on Saturdays and Sundays. We plot the corresponding decline in pre-closure traffic in the right panel. The basic pic-



Notes: NAICS codes 44, 45, 71, 72 only. Only POIs with at least 10 average daily visitors during first week of March and only POIs which are covered all days. POI closure identified on first of first three consecutive days where daily traffic is below 20 percent of average daily traffic during baseline week (first week of March 2020). Closures March 11 and after. The left panel plots the number of daily visits relative to same day-of-week during baseline period (first week of March). The right panel does the same but only for POIs that, in the baseline week, receive at least 50 percent of their total traffic on the weekend.

Figure 9: Declining POI Visits Prior to Closing Date

ture remains the same, although notably visits to “weekend places” was noticeably higher than normal until a few days before they shut down.

B Computational Algorithm

We seek to solve the model with $\alpha \in [0, 1]$, as described in 4.4. Recall that $\alpha = 0$ corresponds to the laissez-faire equilibrium and $\alpha = 1$ corresponds to the planner’s problem. The model is described by the state equations (7)–(8), the optimality condition (18), the costate equations (19)–(20), and the transversality condition (21).

Note that if $u(a) = \log a - a + 1$ or $u(a) = -(1 - a)^2/2$, we can solve equation (18) explicitly for $A(t)$ as a function of the state and costate variables. We use this in our solution and so have a system of four ordinary differential equations.

Our main computational approach numerically solves the state and costate equations using a backward shooting algorithm. We normalize the terminal date to 0 and set $N_i(0) = 10^{-9}$. We then pin down the terminal values of the costate variables by solving equations (19)–(20) with $A(0) = 1$ and $\lambda'_s(0) = \lambda'_i(0) = 0$. Note that the

solution to the differential equation system is not very sensitive to the terminal values of the costate variables, particularly at dates in the distant path.

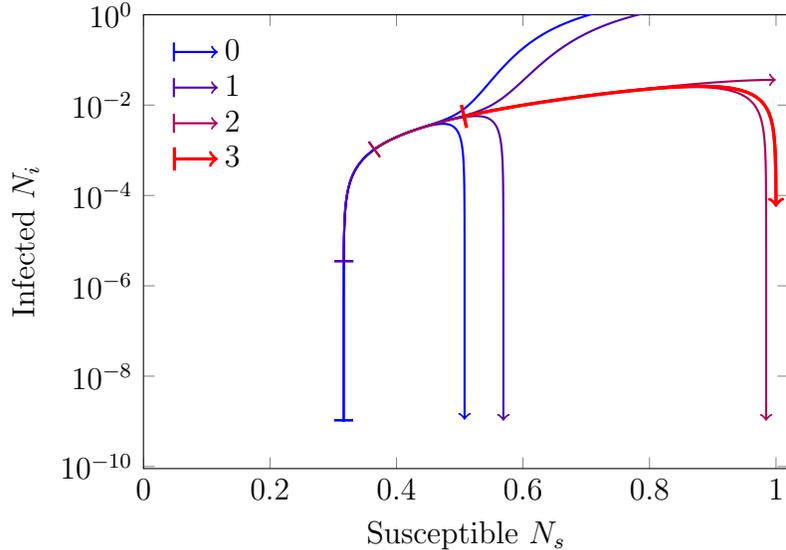
Finally, we search for the terminal value $N_s(0)$. For a given value of $N_s(0)$, we can numerically solve the differential equation system backwards through time. Our aim is to solve backwards to an unknown time $T < 0$ with a known value of $N_s(T)$ and $N_i(T)$, e.g. $N_s(T) = 0.9999223$ and $N_i(T) = 0.0000527$ in our benchmark calibration. If our guess of $N_s(0)$ is too large, we find that there is a finite date t with $N_s(t) = 0.9999223$ and $N_i(t) > 0.0000527$. If our guess of $N_s(0)$ is too small, we find that there is a finite date t with $N_s(t) < 0.9999223$ and $N_i(t) = 0.0000527$. Thus we can easily search for the value of $N_s(0)$ that takes us back to the desired initial condition.

A difficulty arises because the state equations are unstable when we solve backwards in time. Given issues of numerical precision, this means that we may not be able to find a value of $N_s(0)$ for which $N_s(T) = 0.9999223$ and $N_i(T) = 0.0000527$. Instead, two nearby values may lead to widely divergent paths. We illustrate this with the blue lines in Figure 10. The two paths are similar for a while, but ultimately diverge by a substantial amount. The correct value of $N_s(0)$ lies in between the value on these two curves, but we are unable to find it directly.

Instead, we look for the first time that the two paths diverge substantially, defined as the ratio of any of N_s , N_i , λ_s , and λ_i across the two paths exceeding a tolerance threshold, which we set to 10^{-6} . Call the first (i.e. least negative) date when this occurs T_1 . We have now found the solution to the differential equation between T_1 and 0. We then fix the value of N_i , λ_s , and λ_i at the saddle path value at T_1 and again search for the value of $N_s(T_1)$ which takes us back to our desired initial condition, e.g. $N_s(T) = 0.9999223$ and $N_i(T) = 0.0000527$. Once again, we may find two paths near the saddle path which diverge by an amount that exceeds our tolerance. In this case, we call the first (least negative) such date T_2 and repeat the procedure.

Figure 10 plots the joint path of the state variables under laissez-faire and our main calibration. We highlight three of the splits.²² The path labeled ‘1’ starts at $T_1 = -3050$ with $N_s(T_1) = 0.317$ and $N_i(T_1) = 0.00000332$. The path labeled ‘2’ starts at $T_2 = -5110$ with $N_s(T_2) = 0.364$ and $N_i(T_2) = 0.00101$. Finally, the red path label ‘3’ starts at $T_3 = -5520$ with $N_s(T_3) = 0.504$ and $N_i(T_3) = 0.00551$. This

²²In fact, there are five splits in total. The first two splits occur very early on, and we do not plot them.



Notes: The figure corresponds to the laissez-faire version of the model ($\alpha = 0$) in the baseline calibration.

Figure 10: Saddle Path

one reaches $N_s(T) = 0.9999223$ and $N_i(T) = 0.0000527$ at $T = -5786$, our desired terminal condition.

Finally, in plotting the results in the text, we use an alternative normalization of time, defining date T to be March 13, 2020. This does not require resolving the model, just changing the time index.

C Robustness

C.1 Robustness to Parameters and Functional Forms

In this section, we show the laissez-faire and optimal dynamics with reasonable changes to most of the parameters. Given the large uncertainty around many of the key model parameters, fairly large changes are reasonable. We find that our main findings are robust to alternative parameter choices. In particular, we find a strong laissez-faire equilibrium reduction in social activity; an immediate and persistent optimal reduction in social activity that only disappears in the very long run; and an optimal path that keeps the effective reproduction number very close to 1 after peak

Calibration	Laissez-Faire			Optimum		
	Health	Activity	Total	Health	Activity	Total
Baseline	\$11,374	\$1,277	\$12,652	\$860	\$7,227	\$8,087
$u(a) = -\frac{1}{2}(1-a)^2$	\$9,948	\$1,435	\$11,384	\$214	\$5,616	\$5,830
$\kappa = 98.5$	\$6,763	\$611	\$7,373	\$3,942	\$2,358	\$6,300
$R_0 = 6.4$	\$15,973	\$1,799	\$17,772	\$8,776	\$6,630	\$15,406
$\delta = 0.33/365$	\$13,541	\$1,293	\$14,834	\$7,883	\$4,735	\$12,618
$N_i(0) = 0.01$	\$11,745	\$1,327	\$13,072	\$1,852	\$7,164	\$9,016

Notes: Columns 1 and 4: present value of per capita pure health cost of disease. Columns 2 and 5: present value of per capita cost of reduction in social activity. Columns 3 and 6: total cost. First three columns are for the laissez-faire equilibrium, last three are for the optimum. Each row corresponds to a different calibration. We convert utils to US dollars using a factor of 123, as discussed in Section 5.1.

Table 3: Per Capita Welfare Cost of Disease: Robustness

infections.

While the equilibrium amount of social distancing and the health costs of the pandemic are robust to parameter changes, we find some sensitivity under the optimal policy. Table 3 summarizes the total cost under both equilibrium and optimum, along with their breakdown into health and activity cost. When the gains from social distancing are smaller because the health cost of infection κ is lower or the arrival rate of a cure δ is lower, we find that optimal policy allows for an initial wave of infections. The same thing happens when the basic reproduction number R_0 is higher. In this case, effective social distancing requires a bigger reduction in activity, which the planner finds too costly. The resulting wave of infections shifts optimal costs away from reduced activity and towards the health side.

We now discuss each of the modifications in turn, along with its impact on the dynamics of the pandemic.

Alternative Utility Function In Figure 11, we modify the utility function to $u(a) = -\frac{1}{2}(1-a)^2$. This leaves $u(1) = u'(1) = 0$ and also leaves $u''(1)$ unchanged. The latter assumption ensures the value of a statistical life is unchanged. However, it implies less curvature in the utility function and in particular that $u'(0)$ is finite. This modification reduces the marginal value of social activity, and so both equilibrium and

optimal social activity fall with this calibration. Still, all of our takeaway messages hold with this calibration. Optimal policy has an immediate and sustained reduction in social activity, although the extent of it is limited and so the effective reproduction number remains close to 1.

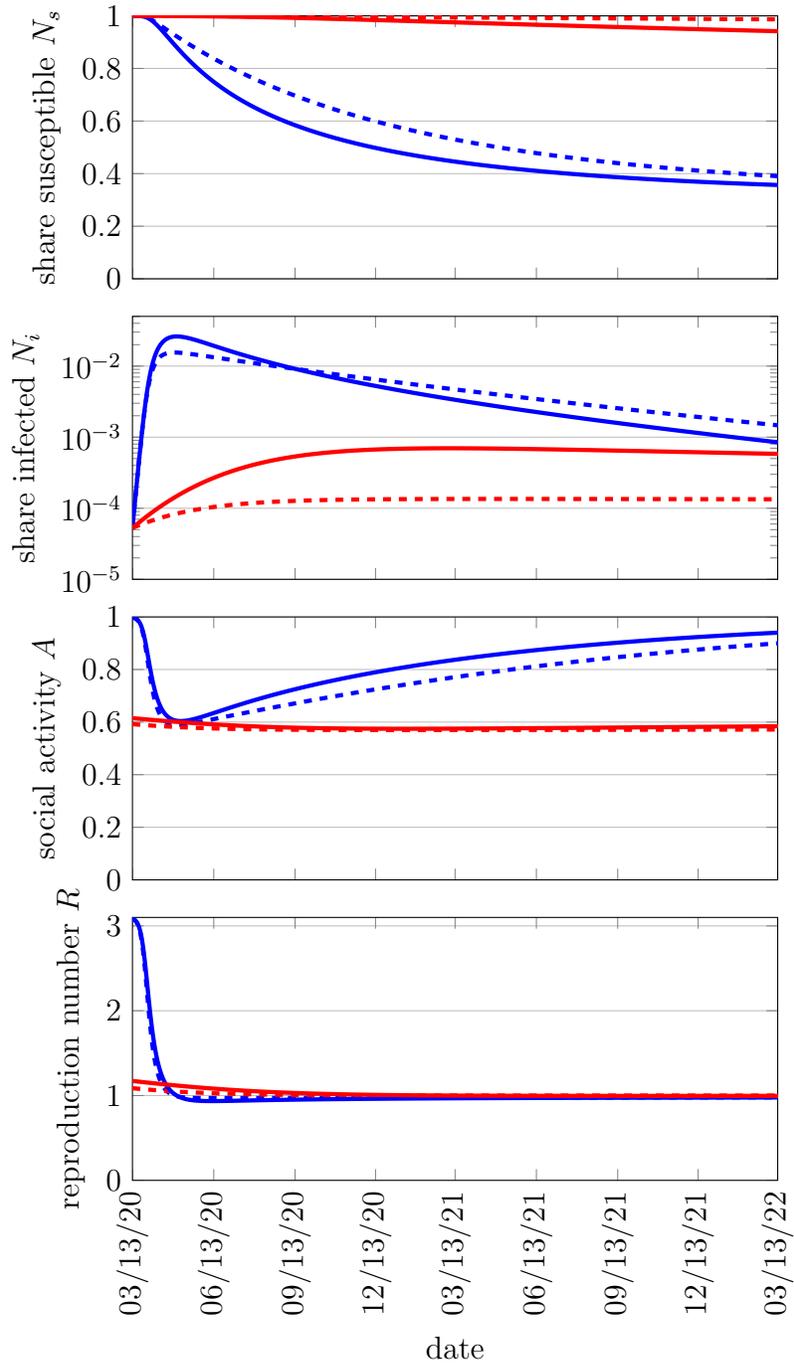
Cost of Disease Here we show the results for $\kappa = 98.5$, that is we cut our baseline parameter value of the expected cost of infection in half. This can be viewed as capturing a lower infection fatality rate or a lower VSL. We show the resulting dynamics for optimum and laissez-faire in Figure 12.

Cutting the cost of infection increases the peak equilibrium infection rate. The peak optimal infection rate increases by a much larger amount, from a very low level of 0.07 percent to a peak close to equilibrium. The delay induced by the optimal policy is also decreased considerably, and the peak of the infection rate under optimal policy is less than 100 days after the equilibrium peak. With a lower VSL, the cost of persistent social distancing exceeds its health care returns.

Higher Duration of Infectivity and Higher R_0 Some authors use a considerably longer duration of infectivity $\frac{1}{\gamma}$. For instance, Hall, Jones and Klenow (2020) set $\gamma = \frac{1}{18}$. We follow them here but maintain the target of a 30 percent daily growth rate in a world without social distancing. We therefore adjust $\beta = 0.3 + \frac{1}{18}$. This gives a substantially higher basic reproduction number of $R_0 = 6.4$.

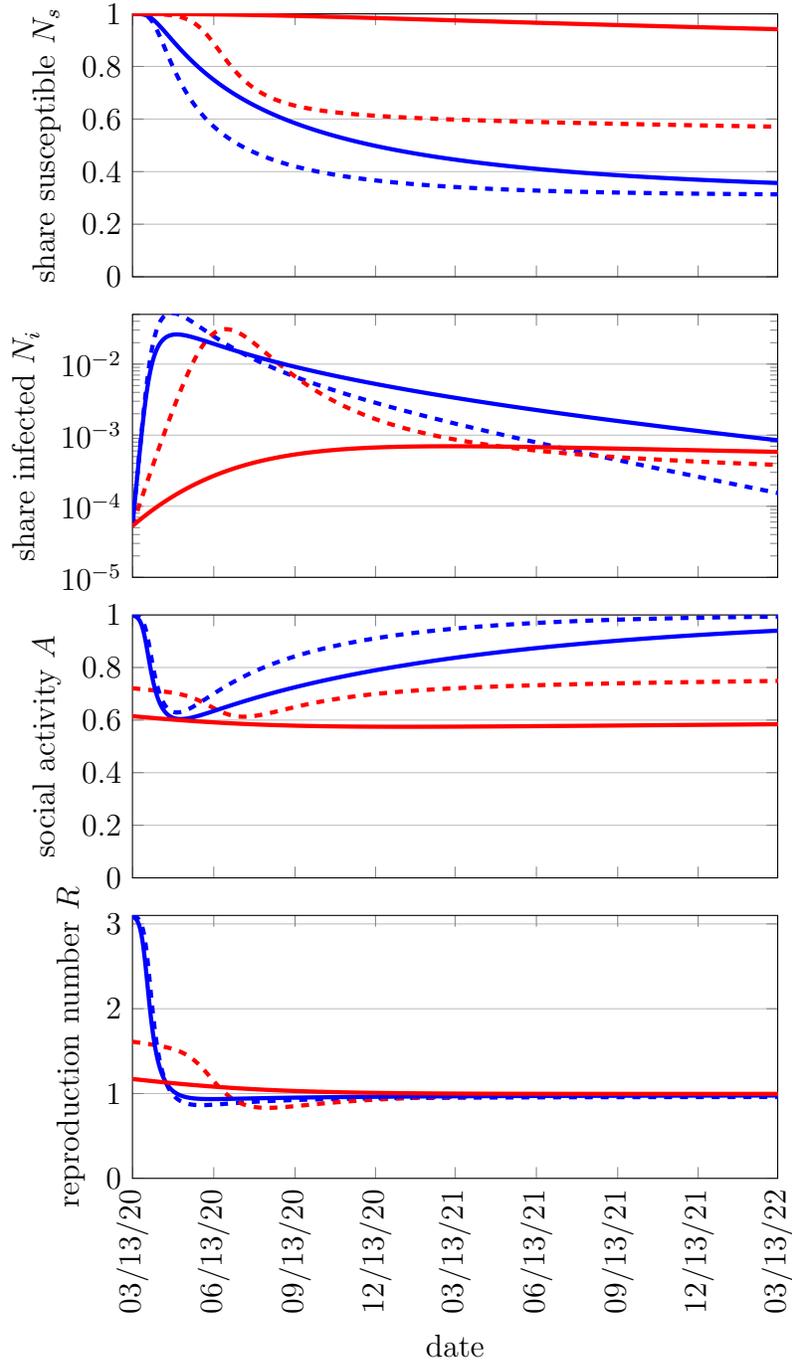
We report the corresponding results in Figure 13. This has qualitatively little impact on the results, but there are significant quantitative differences. Most noticeably, with a much higher basic reproduction number, social distancing is less effective at reducing the infection rate. The peak infection rate in equilibrium is substantially elevated. Even more noticeably, optimal policy now allows for a substantial wave of infections, since the cost of suppressing it is prohibitive. Conversely, both equilibrium and optimal policy see an even larger contraction of social activity. In either case, optimal policy still keeps $R(t)$ marginally below 1 after peak infections.

Lower Cure Probability We have so far assumed that the expected time until a cure is 1.5 years. Figure 14 shows what happens if we double this to 3 years. The laissez-faire equilibrium dynamics are effectively the same, reflecting the fact that



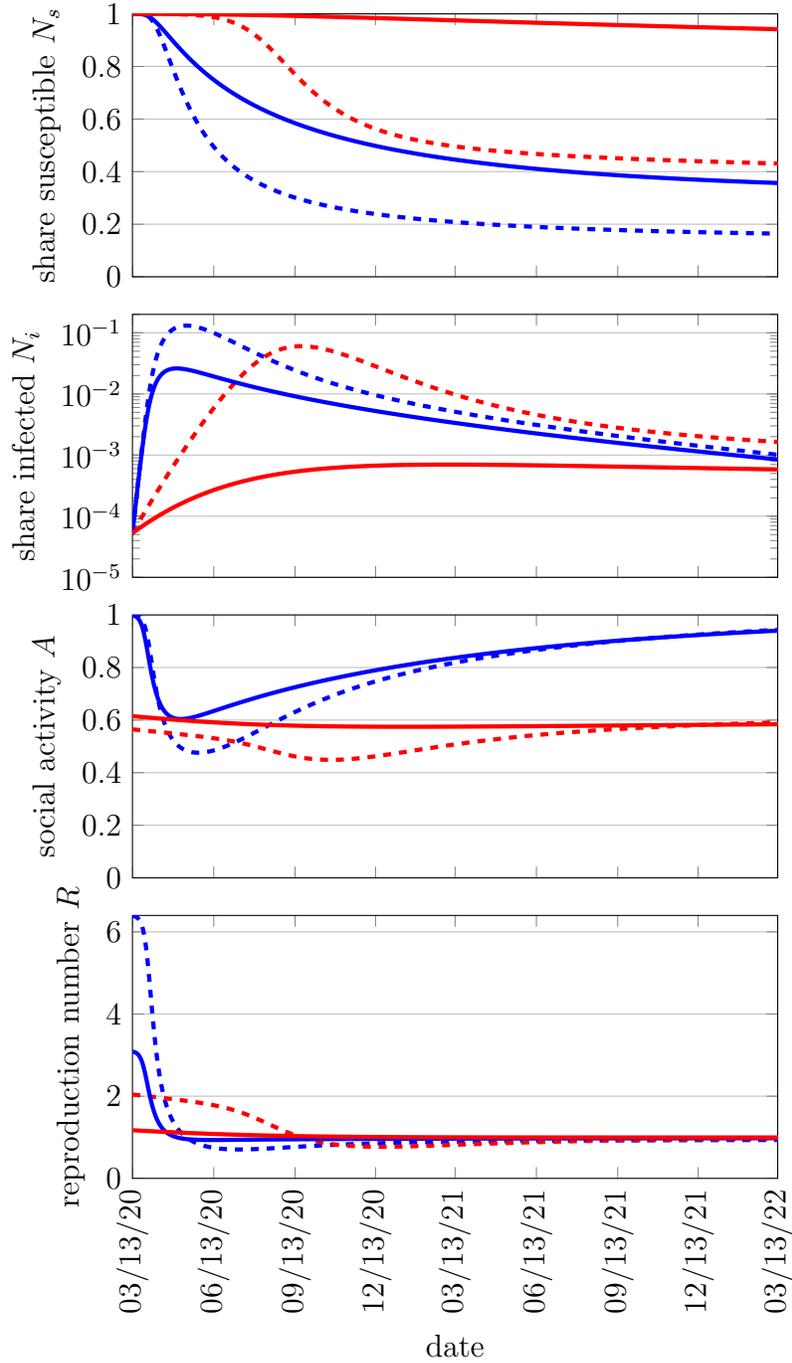
Notes: See Table 1 for calibration. Only change: We assume the utility from social activity is $u(a) = -\frac{1}{2}(1-a)^2$. Solid lines show the baseline calibration and dashed lines show the robustness check. The second plot is drawn on a log scale.

Figure 11: Optimal Policy vs Laissez-Faire with quadratic utility.



Notes: See Table 1 for calibration. Only change: We set $\kappa = 98.5$ instead of $\kappa = 197$. Solid lines show the baseline calibration and dashed lines show the robustness check. The second plot is drawn on a log scale.

Figure 12: **Optimal Policy** vs **Laissez-Faire** with low cost of infection.



Notes: See Table 1 for calibration. Only change: We set $\gamma = 0.056$ (instead of $\gamma = 0.143$) and $\beta = 0.356$ such that $R_0 = 6.4$ (instead of 3.1). Solid lines show the baseline calibration and dashed lines show the robustness check. The second plot is drawn on a log scale.

Figure 13: **Optimal Policy** vs **Laissez-Faire** with long duration of infectivity.

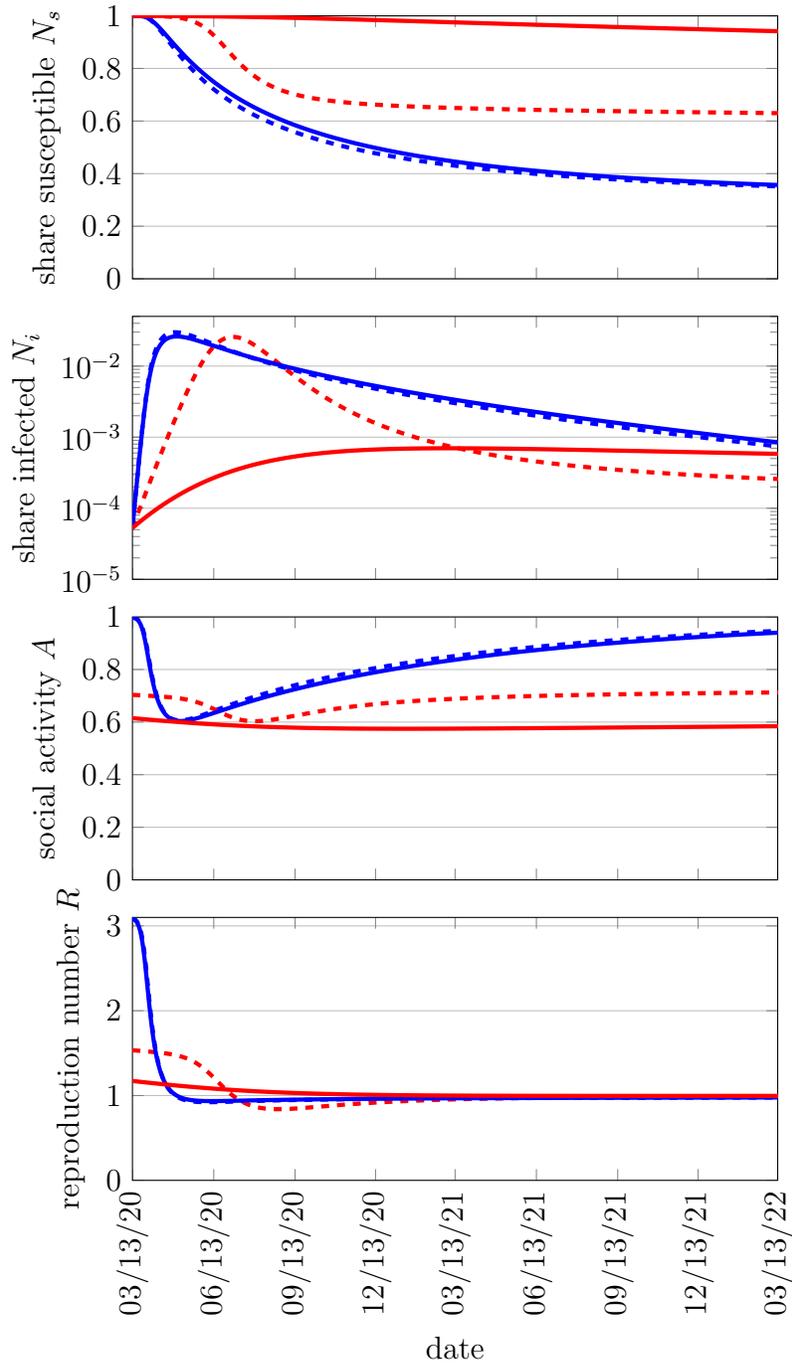
individuals are insensitive to the discount rate, and difficulties in finding a cure are equivalent to a reduction in discounting. The dynamics of the disease, however, differ sharply under the optimal policy. In particular, the planner allows a substantial wave of infections to occur, so the expected number of sick (and hence fatalities) rises much more quickly. This reflects a reduction in the benefits of delay.

Still, the path of optimal policy does not change qualitatively. Immediate and long-lasting social distancing is optimal but the optimal effective reproduction number stays close to 1 after peak infections.

Stock of Initially Infected We have experimented with the stock of the initially infected for several reasons. First, one might think that as the fraction of initially infected individuals becomes exceedingly small one of our key policy lessons—that the optimal policy immediately curtails social activity to buy time—no longer holds. This is indeed true in the limit: for $N_i(0)$ small enough, the expected uncurtailed outbreak date is so far in the future that the (social) gains from social distancing must vanish. However, suppose we start with $N_i(0) = 1/7 \cdot 10^9$, i.e. patient zero, a natural lower bound on the initial seed. With $\beta - \gamma = 0.3$, in just 30 days $N_i(t)$ exceeds 10^{-6} if there is no social distancing. The delay motive we have appealed to above therefore remains quantitatively powerful.

Second, we have also substantially increased $N_i(0)$ up to 1 percent. This could for instance capture the situation in places where COVID-19 appeared first, or it could capture the situation in places where the initial policy response was botched or individuals did not respond because of, say, false information. We plot the corresponding time paths for equilibrium and optimum in Figure 15.

The qualitative patterns are largely unchanged but the optimal policy acts more aggressively. An important observation, shown in the bottom panel of the figure, is that in this case optimal policy suppresses social activity immediately to such an extent that $N_i(t)$ falls from the outset. As a consequence of the high initial stock of infections, the planner has less room to delay the wave of infections and so drives down the stock of infections from the outset. Nonetheless, optimal policy monotonically relaxes social distancing, ultimately relying on a reduction in the share of susceptible individuals for the continued decline in infections. The equilibrium time path, in



Notes: See Table 1 for calibration. The only change is that we set $\delta = \frac{0.33}{365}$ (instead of $\frac{0.67}{365}$). Solid lines show the baseline calibration and dashed lines show the robustness check. The second plot is drawn on a log scale.

Figure 14: **Optimal Policy** vs **Laissez-Faire** with low probability of cure.

turn, remains largely unchanged. Starting with more infections just shifts it ahead in time.

C.2 No Information about Recovery

Throughout the paper, we have assumed that people know when they have recovered from the disease and that recovery confers lifelong immunity. This implies that the recovered's choice of social activity is unaffected by the pandemic, $a_r(t) = A_r(t) = a^* = 1$. Here we change the first part of that assumption. We still assume that recovery confers lifelong immunity, but we assume that infected people pay a cost pay $v = \kappa/\pi$ if they die and otherwise are unaware that they were sick. This means that susceptible, infected, and recovered individuals all have the same information and so must choose the same level of social activity $a(t)$. Similarly, the social planner applies any social distancing requirements to all living individuals.

If individuals do not know their health status, they maximize

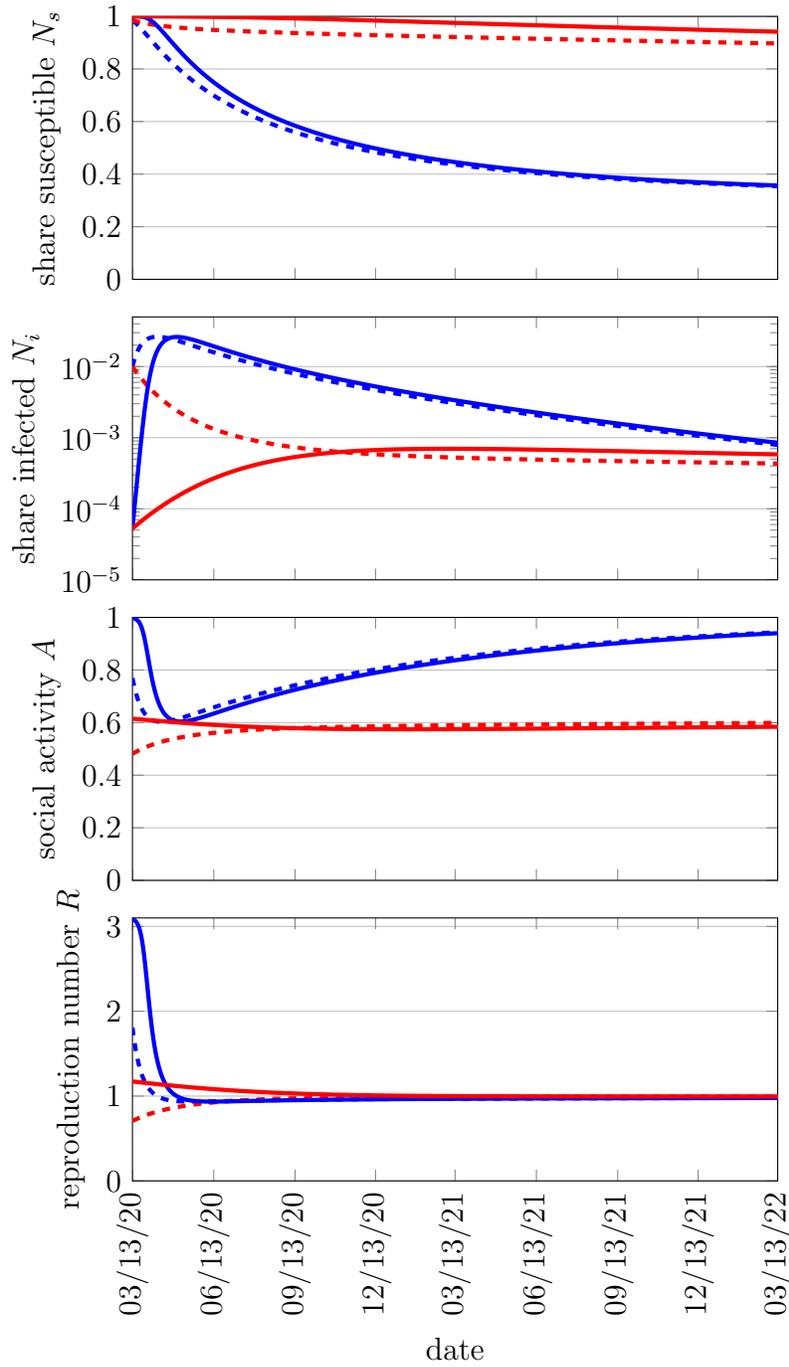
$$\max_{\{a(t)\}} \int_0^{\infty} e^{-(\rho+\delta)t} ((n_s(t) + n_i(t) + n_r(t))u(a(t)) - \gamma n_i(t)\kappa) dt$$

subject to unchanged law of motions for the states, equations (7) and (8). Since a fraction π of the individuals who are not susceptible or infected have died, the remaining fraction $1 - \pi$ are recovered, $n_r(t) = (1 - \pi)(1 - n_s(t) - n_i(t))$. That implies $n_s(t) + n_i(t) + n_r(t) = 1 - \pi(1 - n_s(t) - n_i(t))$. The Hamiltonian can then be written without an additional co-state variable as

$$\begin{aligned} H(n_s(t), n_i(t), a(t), \lambda_s(t), \lambda_i(t)) = & (1 - \pi(1 - n_s(t) - n_i(t)))u(a(t)) - \gamma n_i(t)\kappa \\ & - \lambda_s(t)\beta a(t)n_s(t)A(t)N_i(t) + \lambda_i(t)(\beta a(t)n_s(t)A(t)N_i(t) - \gamma n_i(t)). \end{aligned}$$

The static first order condition is

$$(1 - \pi(1 - n_s(t) - n_i(t)))u'(a(t)) = (\lambda_s(t) - \lambda_i(t))\beta n_s(t)A(t)N_i(t). \quad (23)$$



Notes: See Table 1 for calibration. Only change: We set $N_i(0) = 10^{-2}$ instead of $N_i(0) = 0.0000527$. This also reduces the initial value of $N_s(0)$ to 0.985 instead of 0.9999223. Solid lines show the baseline calibration and dashed lines show the robustness check. The second plot is drawn on a log scale.

Figure 15: **Optimal Policy** vs **Laissez-Faire** with higher stock of initially infected.

The derivatives with respect to the state variables $n_s(t)$ and $n_i(t)$ are

$$(\rho + \delta)\lambda_s(t) - \lambda'_s(t) = \pi u(a(t)) + (\lambda_i(t) - \lambda_s(t))\beta a(t)A(t)N_i(t), \quad (24)$$

$$(\rho + \delta)\lambda_i(t) - \lambda'_i(t) = \pi u(a(t)) - \gamma(\kappa + \lambda_i(t)). \quad (25)$$

The transversality conditions are unchanged

$$\lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_s(t) n_s(t) = \lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \lambda_i(t) n_i(t) = 0. \quad (26)$$

As usual, we impose the equilibrium restrictions $a(t) = A(t)$, $n_s(t) = N_s(t)$, and $n_i(t) = N_i(t)$ to get a set of differential equations that can be solved with the state equations (7) and (8).

The Hamiltonian for the planner is

$$\begin{aligned} H(N_s(t), N_i(t), A(t), \mu_s(t), \mu_i(t)) &= (1 - \pi(1 - N_s(t) - N_i(t))u(A(t)) - \gamma N_i(t)\kappa \\ &\quad - \mu_s(t)\beta A(t)^2 N_s(t)N_i(t) + \mu_i(t)(\beta A(t)^2 N_s(t)N_i(t) - \gamma N_i(t)). \end{aligned}$$

The necessary first order condition with respect to the control A is

$$(1 - \pi(1 - N_s(t) - N_i(t))u'(A(t)) = 2(\mu_s(t) - \mu_i(t))\beta A(t)N_i(t)N_s(t), \quad (27)$$

while the necessary costate equations are

$$(\rho + \delta)\mu_s(t) - \mu'_s(t) = \pi u(A(t)) + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_i(t), \quad (28)$$

$$(\rho + \delta)\mu_i(t) - \mu'_i(t) = \pi u(A(t)) + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_s(t) - \gamma(\kappa + \mu_i(t)). \quad (29)$$

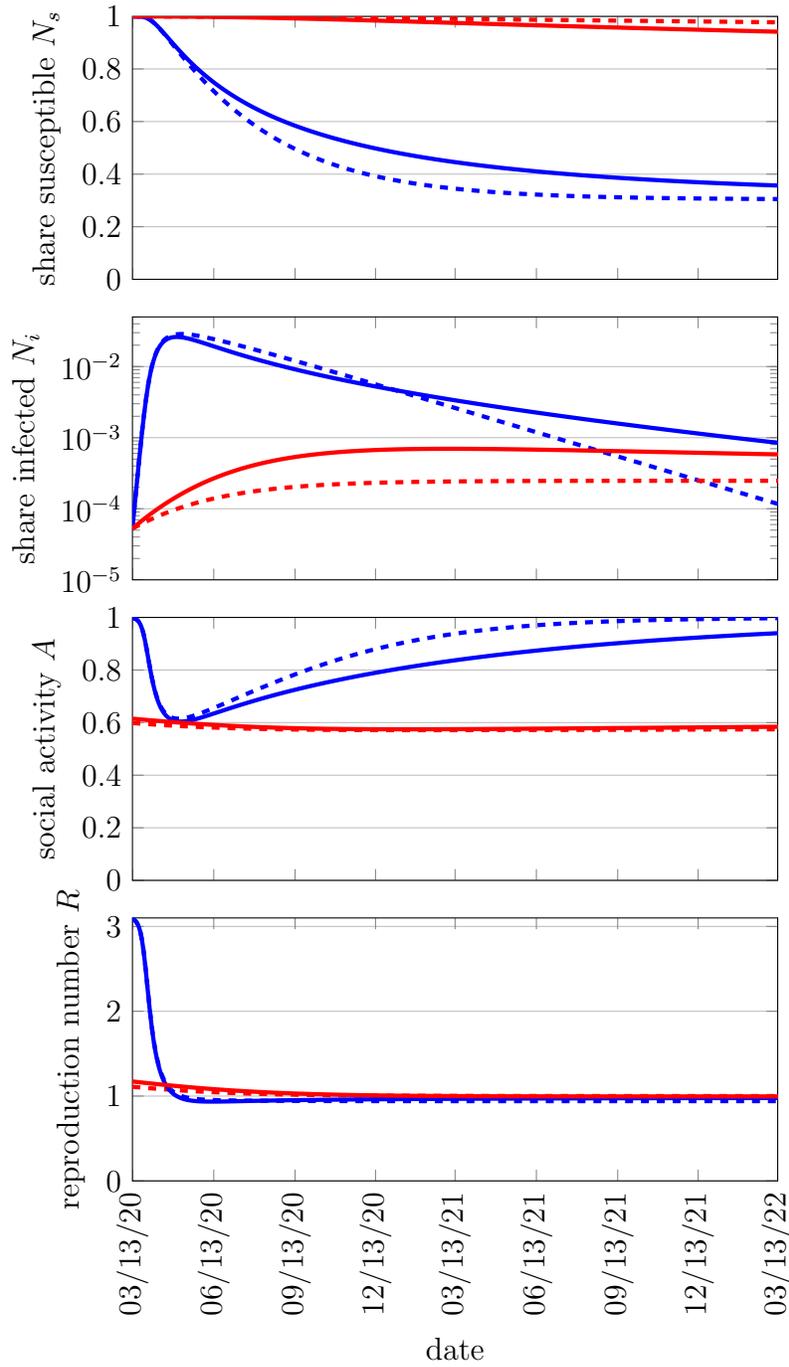
Finally, the planner also has unchanged necessary transversality conditions

$$\lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \mu_s(t) N_s(t) = \lim_{t \rightarrow \infty} e^{-(\rho+\delta)t} \mu_i(t) N_i(t) = 0. \quad (30)$$

Again, we solve these differential equations with the state equations (7) and (8).

Figure 16 shows the results. If individuals do not know when they have recovered, they increase their social activity under laissez-faire when susceptible or infected, compared with the benchmark model. This reflects the fact that as the disease progresses,

they are increasingly likely to have acquired immunity. In contrast, the optimal policy suppresses activity when the planner is unaware of recovery status. This reflects the opposing force, that acquiring immunity is less valuable when people who are immune must also be subject to social distancing requirements.



Notes: See Table 1 for calibration. Only change: We assume individuals do not know their recovery status. Solid lines show the baseline assumption and dashed lines show the robustness check. The second plot is drawn on a log scale.

Figure 16: **Optimal Policy** vs **Laissez-Faire** when individuals do not know whether they have recovered.