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

We thank SafeGraph for making their data freely available to the research community. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 JEL No. E1,H0,I1

ABSTRACT

We use a conventional dynamic economic model to integrate individual optimization, equilibrium interactions, and policy analysis into the canonical epidemiological model. Our tractable framework allows us to represent both equilibrium and optimal allocations as a set of differential equations that can jointly be solved with the epidemiological model in a unified fashion. Quantitatively, the laissez-faire equilibrium accounts for the decline in social activity we measure in US micro-data from SafeGraph. Relative to that, we highlight three key features of the optimal policy: it imposes immediate, discontinuous social distancing; it keeps social distancing in place for a long time or until treatment is found; and it is never extremely restrictive, keeping the effective reproduction number mildly above the share of the population susceptible to the disease.

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

Gregor Jarosch Department of Economics Princeton University 195 Julis Rabinowitz Building Princeton, NJ 08544 and NBER gregorjarosch@gmail.com Robert Shimer Department of Economics University of Chicago 1126 East 59th Street Chicago, IL 60637 and NBER shimer@uchicago.edu

1 Introduction

A key parameter in workhorse models of disease transmission is the rate at which sick people infect susceptible people. A large set of policy measures such as compulsory social distancing aim at reducing this rate. In this paper, we model the rate of transmission as reflecting the choice of rational individuals who weigh the cost of getting infected against the benefits derived from social activity. These benefits capture both social and economic returns from physical human interaction.

To motivate the exercise, we use micro-data from SafeGraph to show that individuals across the United States substantially reduced their exposure to others long before state and local governments imposed the first "shelter-in-place" restrictions in response to the Covid-19 pandemic. We then show how to integrate such optimizing behavior into an otherwise standard epidemiological model. Specifically, we use optimal control theory to derive two ordinary differential equations (ODEs) which capture individual optimality. Together with two standard differential equations from epidemiological models, the resulting system of four differential equations fully summarizes the model and can easily be solved.

The model is consistent with the observation that social activity fell drastically before there were any legally-mandated restrictions on movement. It is thus a natural laboratory for evaluating social distancing policies. To do so, we also show how to characterize the symmetric Pareto optimal allocation. Optimal policy chooses a time path for the amount of social activity, recognizing the health consequences of a high level of activity. Optimal policy can likewise be described by the solution to a simple system of four ODEs. Comparison with the laissez-faire benchmark elucidates the external effects in disease transmission. Moreover, we show how perfect altruism eliminates the gap between equilibrium and optimum.

The internal benefits of social distancing come from the fact that someone is less likely to get sick if they engage in more social distancing. This is reflected in individual behavior. The external benefit comes from the fact that they are less likely to get other people sick, particularly other strangers. While individuals internalize the cost of social distancing, optimal policy also recognizes that individuals may ignore the external benefit of reducing the risk of transmitting illness. Moreover, optimal policy internalizes the effect of an additional sick person on the quality of health care available to inframarginal individuals.

We then turn quantitative. Since the model is very parsimonious, the calibration targets various epidemiological findings such as the initial growth rate of the disease, the duration of infectivity, and the fatality rate.

Our most important findings are the following: First, the laissez-faire equilibrium reduction in social activity due to the internalized cost of infection is strong. It delays the peak outbreak and leads to a substantial reduction in the number of expected fatalities, relative to a no-response benchmark. However, individuals only reduce activity once the risk of infection becomes non-negligible.

Second, and in contrast, social distancing optimally starts as soon as the disease emerges, discontinuously suppressing social activity. This discrete drop in activity delays the pandemic and hence buys time. Because of the hope for a cure, this strictly reduces the expected number of deaths and yields a welfare gain.

Third, optimal social distancing is persistent. Absent a cure, social activity remains depressed for years. This is the flip side of delay. Because the pool of susceptible individuals drains only slowly, activity needs to remain suppressed for a long time unless a cure is found. Nevertheless, asymptotically social activity returns to normal, with infections stopping only because of herd immunity as the product of the disease's basic reproduction number R_0 and the share of people who are susceptible falls below one.

Fourth, social distancing is never too restrictive. At any point in time, the *effective* reproduction number for a disease is the expected number of people that an infected person infects. In contrast to the basic reproduction number, it accounts for the current level of social activity and the fraction of people who are susceptible. Importantly, optimal policy keeps the effective reproduction number above the fraction of people who are susceptible, although for a long time only mildly so. That is, social activity is such that, if almost everyone were susceptible to the disease, the disease would grow over time. That means that optimal social activity lets infections grow until the susceptible population is sufficiently small that the number of infected people starts to shrink. As the stock of infected individuals falls, the optimal ratio of the effective reproduction number to the fraction of susceptible people grows until it eventually converges to the basic reproduction number.

To understand why social distancing is never too restrictive, first observe that social activity optimally returns to its pre-pandemic level in the long run, even if a cure is never found. To understand why, suppose to the contrary that social distancing is permanently imposed, suppressing social activity below the first-best (disease-free world) level. That means that a small increase in social activity has a first-order impact on welfare. Of course, there is a cost to increasing social activity: it will lead to an increase in infections. However, since the number of infected people must converge to zero in the long run, by waiting long enough to increase social activity, the number of additional infections can be made arbitrarily small while the benefit from a marginal increase in social activity remains positive.

Now consider the role of social distancing in the short run. With a low initial infection rate, pushing the effective reproduction number below the share of susceptible people implies that the total number of individuals who get infected over any time interval will be small. That means that the health status of the population—the share of susceptible and infected people—will barely change. It follows that if it is optimal to keep the effective reproduction number below the share of susceptible people initially, it will be optimal to do so at any later date. But we have just argued that this cannot be the case. Thus the effective reproduction number must always be bigger than the share of susceptible people. In our calibrated model, it is mildly so for a sustained period. We note that this argument is predicated on the assumption that the initial infection rate is small, since otherwise a period of strong social distancing may have a big effect on the health status of the population. Indeed, we verify that if the initial infection rate is large enough, an optimal policy may temporarily push the effective reproduction number below the share of susceptible people.

We then revisit our micro-data and contrast it quantitatively with the model under our baseline calibration. The model captures both the drop in social activity prior to any government intervention and the pace of the contraction surprisingly well. We corroborate this with additional aggregated data from Sweden that display similar patterns.

We then consider several robustness exercises with respect to our calibration. An important takeaway is that even large parameter changes matter little for the shape of equilibrium or optimal social activity. Optimal policy is fairly insensitive even to large parameter changes and is well summarized by the points discussed above: Immediate social distancing that ends only slowly but is not overly restrictive. This is reassuring given the large current amount of parameter uncertainty. These robustness exercises also document that our basic framework naturally accommodates a rich set of extensions. We therefore conclude with an additional set of proposed extension that we believe are of first order.

2 Related Work

Our basic approach builds on the susceptible-infected-recovered (SIR) model (Kermack and McKendrick, 1927). There is a rapidly growing body of work that uses this epidemiological model, together with standard economic models, to understand the interplay between disease transmission and economic activity.

The basic epidemiology model is reviewed in Atkeson (2020), who analyzes the optimal lock down policy in an economic environment. Following a similar approach, Eichenbaum, Rebelo and Trabandt (2020) study the two-way interaction of disease dynamics and economic activity in a macroeconomic SIR model. While theirs is a substantially richer environment, we obtain a larger degree of analytical tractability and treat equilibrium and optimum in a unified fashion that simply adds two ODEs to the SIR model. A further key difference is that, in their model, disease transmission and its externalities happens through consumption and the government acts through a tax on consumption. In our setup, disease transmission happens through general activity, both economic and social and the government acts through restricting social activity. This allows us to directly map to newly available data on social activity, for example the SafeGraph data on foot traffic we discuss below.

Other papers that focus on the individual response to a pandemic include Garibaldi, Moen and Pissarides (2020), who use tools from search and matching and Krueger, Uhlig and Xie (2020), who focus on the shift in the sectoral composition of consumption as a force that mitigates the economic fallout of the pandemic.

Alvarez, Argente and Lippi (2020) study a planning problem similar to ours where the planner directly controls the amount of activity and trades of the losses from restrictions against the health benefits. However, they exogenously fix the amount of activity absent policy intervention, while we explicitly model the choice of social activity by individuals. This allows us to contrast the optimal amount of social contacts with the laissez-faire one, taking seriously that even absent mobility restrictions the negative effects of the epidemic are partially internalized.

Some other recent papers focusing on equilibrium and optimal policy include Jones, Philippon and Venkateswaran (2020), Hall, Jones and Klenow (2020), Dewatripont, Goldman, Muraille and Platteau (2020), Piguillem and Shi (2020), Barro, Ursua and Weng (2020), Glover, Heathcote, Krueger and Rios-Rull (2020), Keppo, Kudlyak, Quercioli, Smith and Wilson (2020), and Kaplan, Moll and Violante (2020).

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. E.g., with more individuals in parks, restaurants and public transit any given trip to a park/restaurant/subway visit is more likely to lead to disease. Such a quadratic matching function has a search externality that, traditionally, is viewed as positive (Diamond, 1982), but that turns negative in an age of disease. It stands in contrast to a linear search technology where an individual's 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 activity and the associated risk of disease transmission is 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.

Greenwood, Kircher, Santos and Tertilt (2019) study the HIV epidemic in Malawi using a computational choice-theoretic equilibrium model of sexual behavior. They model the individual effort to find a partner in different markets, which are associated with different degrees of risk. And indeed, the risk of infection at any given market depends on the entire distribution of health types visiting the market in line with the arguments just made.

Budish (2020) treats disease containment as an economic constraint and discusses policies that maximize welfare subject to the containment constraint. In turn, our framework treats the system of differential equations governing the evolution of the sickness as the relevant constraint. A policy-maker maximizes welfare subject to that constraint fully taking into account the damage caused by the disease. As a consequence, the policy-maker may well choose policies that let the epidemic spread if the cost of containment are too high. Indeed, we find that this is always optimal, albeit at a slow pace.

3 Declining Social Activity, Early and Everywhere

In this section, we use newly available micro-data that document substantial behavioral changes across the United States even before any policy measures were taken.

We work with micro-data data from SafeGraph.¹ Among other things, SafeGraph provides highly disaggregated and detailed high-frequency information on individual travel in the United States. The population sample is a panel of opt-in, anonymized smartphone devices, and is well balanced across US demographics and space.

In early April 2020, SafeGraph made two datasets freely available to researchers.² Their first "Covid-19 Response Dataset," named "Weekly Patterns," registers GPS-identified visits to Points of Interest (POI) (primarily businesses) with exact known location in the United States at hourly frequency in a balanced panel. The data is currently available covering the period March 1 to April 11, 2020. The dataset is large. On March 1, the dataset recorded approximately 32.1 million individual visits to approximately 3.9 million POI.

The second dataset, named "Social Distancing Metrics," uses information from individual cell devices that can be assigned to a home address (using their night-time location) to measure individual foot traffic and its response to the outbreak. The dataset goes back to January 1, 2020 and currently, runs until April 9 and is likewise 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, the data measures for each census

¹Attribution: 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.

 $^{^{2}}$ For detailed information, see https://docs.safegraph.com/docs/weekly-patterns and https://docs.safegraph.com/docs/social-distancing-metrics.

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-of-week at the home location (variable completely_home_device_count).

We construct our measures at the state level. We use the first dataset to count the total daily number of visits, for each state, to POIs. We proceed identically for our other two measures. We subtract the median minutes spent at home from $24 \times 60 = 1440$ and take a daily state-wide average. We similarly construct the state-wide fraction of all devices that leave the house at least once during any day. We express all three variables relative to a baseline week (dividing by the corresponding day during the first week of March).

This gives us, for each state, three different measures of the decline of social activity that naturally map to the model. Figure 1 reports our result for the first variable, visits to POI. We plot, at any given date, the median value (across states) of the decline relative to baseline, along with the max and min and 10th and 90th percentile.

The figure shows a remarkably uniform contraction of social activity across US states beginning in the second week of March, leveling off at some 50 percent relative to baseline towards the end of the month.

The figure also depicts the fraction of the US population subject to official stay-at-home or shelter-in-place orders. The figure shows that social activity began contracting 10 days before the first significant orders were put into place around March 20.

We complement this with the other two social distancing metrics we have available, days spent entirely at home and daily dwell time at home. Their decline is depicted in Figure 2. These two variables display a somewhat smaller decline of 20-30% relative to baseline. The basic pattern remains the same: Social activity starting contracting substantially, rapidly, and long before the first lockdown measures. This also happened across the board in the United States.

This offers a direct and readily available measure of the extent of individual social activity which is both the key endogenous variable in our model, as well as 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. We will later 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 the US in March 2020.

4 Model

The basic epidemiological framework is a continuous time SIR model with a possibility of death, i.e. a SIRD model. Individuals susceptible to the disease may become infected



Notes: Visits to Points of Interest in SafeGraph's "Weekly Patterns" Covid-19 Response Dataset. We sum daily visits within each state and express them relative to the same day during baseline week (first week of March 2020). We plot the median (min, max, ... across states) decline relative to baseline at any given day. The solid blue line is the median. The dark share is 10% - 90% interval, and the light shade is min-max interval. The solid red line is the percent of the population subject to stay-at-home or shelter-in-place orders. 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.

Figure 1: Declining Activity, Early and Everywhere



Notes: Social activity based on SafeGraph's "Social Distancing Metrics" Covid-19 Response Dataset. Left Panel: Fraction of devices that leave assigned "home" at least once during any day. Right Panel: Dwell time at home, median device. Measures at the census block group, we take state-wide weighted averages and express them relative to the same day during baseline week (first week of March 2020). We plot the median (min, max, .. across states) decline relative to baseline at any given day. The solid blue line is the median. The dark share is 10% - 90% interval, and the light shade is min-max interval. The solid red line is the population under lockdown. Population under lockdown is based on authors' own calculations using https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html.

Figure 2: Declining Activity, Early and Everywhere II

through contact with other infected individuals. The infected stochastically recover or die. Individuals do not know if they are susceptible or infected, but in our baseline model, they can tell when they have recovered from the disease. In our baseline model, individuals are otherwise homogeneous.

At each time $t \ge 0$, a measure 1 of individuals are in one of the four states, susceptible (s), infected (i), recovered (r), or deceased (d). Let $N_j(t)$, $j \in \{s, i, r, d\}$, denote the measure of individuals in each state. Thus $N_s(t)$ also gives the fraction of the population that has not gotten infected. Assume that $N_i(0) > 0$, so there is a seed of infection.

A susceptible individual can get infected by meeting an infected individual. However, individuals are unaware if they are infected. Infected individuals recover at rate $(1 - \pi(N_i(t)))\gamma$ and die at rate $\pi(N_i(t))\gamma$, where $\pi(N_i(t)) \in [0, 1]$ is the infected fatality rate, the fraction of infected individuals who eventually die from the disease. We allow this to depend on the number of infected people, reflecting the possibility that the disease overwhelms the hospital system. A recovered individual knows that he is recovered. 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 lump the risk and cost of death (i.e. the lost value of life) and the cost of disease together in a single function $\kappa(N_i(t))$, the expected cost that an infected individual pays when he exits the infected state. Again, we allow for this cost to depend on N_i to capture that treatment quality (and hence π) might deteriorate when the health care system becomes overrun.

Individuals discount the future at rate ρ , and 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 the disease.

We assume that all living individuals choose their level of *social activity* a, and get utility u(a) from social activity a. Assume u is single-peaked and its maximum is attained at a finite value $a^* > 0$. 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 the outbreak. The latter allows us to use u as a measure of the utility loss from social distancing.

Disease transmission depends on the rate at which individuals have social interactions. Let $A_j(t)$ denote the aggregate amount of social activity by all individuals of type $j \in \{s, i, r\}$. 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 that an individual of type j has social interaction with an individual of type j' is $A_j(t)N_j(t)A_{j'}(t)N_{j'}(t)$, the product of the level of social activity of the two types. In particular, the rate that susceptible individuals get sick 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.³

Together the assumptions that preferences u depend on social activity while disease transmission depends on social interactions, is central to our view of social distancing. It 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.⁴ 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.

The disease transmission function captures the idea that if a particular group chooses little social activity $A_j(t)$, one is unlikely to make social contact with them. It also captures the idea that the amount of social interactions depends not only on a particular group's choice, but also on everyone else choices. This captures, for instance, that even an individual frequently going to a restaurant or to the office has few social interactions if nobody else is there. As a consequence, disease transmission displays a negative externality: increasing social activity increases other people's social interactions, putting them at a higher risk of infection.

Under our modeling assumptions, the laws of motion describing the aggregate state are given by

$$N'_{s}(t) = -\beta A_{s}(t)N_{s}(t)A_{i}(t)N_{i}(t)$$

$$\tag{1}$$

$$N'_{i}(t) = \beta A_{s}(t) N_{s}(t) A_{i}(t) N_{i}(t) - \gamma N_{i}(t)$$

$$\tag{2}$$

$$N'_{r}(t) = (1 - \pi(N_{i}(t)))\gamma N_{i}(t)$$
(3)

$$N'_d(t) = \pi(N_i(t))\gamma N_i(t) \tag{4}$$

with each of the $N_j(0)$ given. If $A_j(t) = 1$ for $j \in \{s, i\}$ and $\pi(\cdot) = 0$, this is the standard SIR model (Kermack and McKendrick, 1927). Allowing $\pi \neq 0$ gives us the SIRD model. We stress that there are not only social interactions between the susceptible and the infected,

³An alternative would be a linear technology, in which case type j individuals interact with type j' individuals at rate $\frac{A_j(t)N_j(t)A'_j(t)N'_j(t)}{\sum_{j''}A''_j(t)N''_j(t)}$. This formulation makes sense if type j individuals desire to interact with somebody at rate $A_j(t)$. 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). 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 instance the optimal social activity of the recovered, which one would want to boost with a linear matching technology.

⁴For 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.

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 social matching function allows these interactions to happen, but they do not affect the number of interactions where the disease gets transmitted.

Note that only $A_s(t)$ and $A_i(t)$ affect disease transmission. Also note our assumption that individuals do not know whether they are susceptible or infected. We thus impose the measurability restriction that $A_s(t) = A_i(t)$ and use A(t) to denote this common level of social activity. The expected number of others in contact with an infected individual during the time she is infected is $\beta A(t)^2/\gamma$. In a population in which individuals do not change their behavior in response to the disease $A(t) = a^* = 1$, e.g. when the disease first emerged. Therefore the *basic reproduction number* R_0 , defined as the expected number of others infected by an infected individual in a population where everyone is susceptible, is $R_0 = \beta/\gamma$. We also define the *effective reproduction number*, the expected number of others infected by an infected individual, given the current level of social activity A(t) and the current fraction of susceptible people $N_s(t)$. This number is $R_e(t) = \beta A(t)^2 N_s(t)/\gamma$. Reductions in social activity drive the ratio of the effective reproduction number R_0 , which is a primitive of the environment.

5 Laissez-faire Equilibrium

In this section, we consider the problem of an individual who is either susceptible or infected choosing his own rate of social activity, taking the number and social activity of other infected people 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

$$\max_{\{a(t),a_r(t)\}} \int_0^\infty e^{-(\rho+\delta)t} \big((n_s(t) + n_i(t))u(a(t)) + n_r(t)u(a_r(t)) - \gamma n_i(t)\kappa(N_i(t)) \big) dt$$
(5)

subject to

$$n'_{s}(t) = -\beta a(t)n_{s}(t)A(t)N_{i}(t)$$
$$n'_{i}(t) = \beta a(t)n_{s}(t)A(t)N_{i}(t) - \gamma n_{i}(t)$$
$$n'_{r}(t) = (1 - \pi(N_{i}(t)))\gamma n_{i}(t)$$

with $n_s(0) = N_s(0)$, $n_i(0) = N_i(0)$, and $n_r(0) = N_r(0)$ given. The laws of motions reflect that the individual takes the time path of everyone else's choice A(t) and the associated aggregate infection rate $N_i(t)$ as given. However, the individual's past choices of a(t) affects his probability of being in each of the states.

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 then $u(a_r(t)) = 0$ for all t. Dropping this control variable and the unnecessary third constraint, write the current value Hamiltonian as

$$H(n_s(t), n_i(t), a(t), \lambda_s(t), \lambda_i(t)) = (n_s(t) + n_i(t))u(a(t)) - \gamma n_i(t)\kappa(N_i(t)) - \lambda_s(t)\beta a(t)n_s(t)A(t)N_i(t) + \lambda_i(t)(\beta a(t)n_s(t)A(t)N_i(t) - \gamma n_i(t)),$$

where $\lambda_s(t)$ and $\lambda_i(t)$ are the co-state variables 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:

$$(n_s(t) + n_i(t))u'(a(t)) = (\lambda_s(t) - \lambda_i(t))\beta n_s(t)A(t)N_i(t).$$
(6)

This static first order condition balances the returns from social activity and the risk of getting infected. Second and third, 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:

$$(\rho + \delta)\lambda_s(t) - \lambda'_s(t) = u(a(t)) + (\lambda_i(t) - \lambda_s(t))\beta a(t)A(t)N_i(t),$$
(7)

$$(\rho + \delta)\lambda_i(t) - \lambda'_i(t) = u(a(t)) - \gamma \big(\kappa(N_i(t)) + \lambda_i(t)\big).$$
(8)

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

$$\lim_{t \to \infty} e^{-(\rho+\delta)t} \lambda_s(t) n_s(t) = \lim_{t \to \infty} e^{-(\rho+\delta)t} \lambda_i(t) n_i(t) = 0.$$
(9)

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 gives us a system of four differential equations and one static equation. Together with initial conditions for $N_s(0)$ and $N_i(0)$ and the transversality conditions (9), these fully summarize the model:

$$N'_{s}(t) = -\beta A(t)^{2} N_{s}(t) N_{i}(t)$$

$$N'_{i}(t) = \beta A(t)^{2} N_{s}(t) N_{i}(t) - \gamma N_{i}(t)$$

$$(\rho + \delta) \lambda_{s}(t) - \lambda'_{s}(t) = u(A(t)) + (\lambda_{i}(t) - \lambda_{s}(t)) \beta A(t)^{2} N_{i}(t),$$

$$(\rho + \delta) \lambda_{i}(t) - \lambda'_{i}(t) = u(A(t)) - \gamma \big(\kappa(N_{i}(t)) + \lambda_{i}(t)\big),$$

$$(N_{s}(t) + N_{i}(t)) u'(A(t)) = (\lambda_{s}(t) - \lambda_{i}(t)) \beta A(t) N_{s}(t) N_{i}(t).$$

The first two differential equations impose $A_s(t) = A_i(t) = A(t)$ on the aggregate relationships (1) and (2). The last three equations correspond to equations (7), (8), and (6) with $a(t) = A(t), n_i(t) = N_i(t)$, and $n_s(t) = N_s(t)$.

We solve this model through a backward shooting algorithm. Fix $\lambda_s(T)$ and $\lambda_i(T)$ at their asymptotic values at some faraway date T, but make $N_i(T)$ slightly positive. Then search for the value of $N_s(T)$ that achieves a desired initial condition for N_s and N_i at a much earlier date. In practice, we can find the equilibrium in a few seconds.

6 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 illness. We show that this gives rise to a system of four ODEs which very closely resemble the ODEs shaping equilibrium we just discussed.

6.1 Planner's Problem

The planner solves

$$\max_{\{A(t),A_r(t)\}} \int_0^\infty e^{-(\rho+\delta)t} \big((N_s(t) + N_i(t))u(A(t)) + N_r(t)u(A_r(t)) - \gamma N_i(t)\kappa(N_i(t)) \big) dt \quad (10)$$

subject to equations (1)–(3). 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. Then

the Hamiltonian is

$$H(N_s(t), N_i(t), A(t), A_s(t), A_i(t)) = (N_s(t) + N_i(t))u(A(t)) - \gamma N_i(t)\kappa(N_i(t)) - \mu_s(t)\beta A(t)^2 N_s(t)N_i(t) + \mu_i(t) (\beta A(t)^2 N_s(t)N_i(t) - \gamma N_i(t)).$$

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

$$(N_s(t) + N_i(t))u'(A(t)) = 2(\mu_s(t) - \mu_i(t))\beta A(t)N_i(t)N_s(t),$$
(11)

while the necessary costate equations are

$$(\rho + \delta)\mu_s(t) - \mu'_s(t) = u(A(t)) + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_i(t),$$
(12)

$$(\rho + \delta)\mu_i(t) - \mu'_i(t) = u(A(t)) - \gamma \big(\kappa(N_i(t)) + N_i(t)\kappa'(N_i(t)) + \mu_i(t)\big) + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_s(t).$$
(13)

Finally, the planner also has necessary tranversality conditions

$$\lim_{t \to \infty} e^{-(\rho+\delta)t} \mu_s(t) N_s(t) = \lim_{t \to \infty} e^{-(\rho+\delta)t} \mu_i(t) N_i(t) = 0.$$
(14)

There are a few key differences between the first order conditions in the equilibrium and optimum problem. 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 (11) compared to equation (6). Second, the planner recognizes the health care externality, that the cost of being sick may depend on how many people are sick $N_i(t)$. This is the additional term involving the derivative of the cost function κ in equation (13) compared to equation (8). Third, the planner recognizes that sick people get other people sick, while in equilibrium individuals do not care about this outcome. This shows up as the last term in equation (13).

As usual, the aggregate state still satisfies

$$N'_{s}(t) = -\beta A(t)^{2} N_{s}(t) N_{i}(t)$$
$$N'_{i}(t) = \beta A(t)^{2} N_{s}(t) N_{i}(t) - \gamma N_{i}(t)$$

Thus 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 $N_i(T)$.

Before we put numbers into the model, we briefly discuss how a model with altruism can

encompass both laissez-faire and optimum.

6.2 Perfect and Imperfect Altruism

In the laissez-faire equilibrium, 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 capture this through 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, except for any possible congestion in the health care system. We show that this is the social optimum when κ is constant.

To illustrate this, we modify the individual objective function to assume individual are concerned about making others sick:

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

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 $\lambda_i(t) - \lambda_s(t)$, where again $\lambda_j(t)$ is the costate variable on $n_j(t)$, $j \in \{s, i\}$. In words, 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:

$$(n_s(t) + n_i(t))u'(a(t)) = \beta A(t) (\lambda_s(t) - \lambda_i(t)) (n_s(t)N_i(t) + \alpha n_i(t)N_s(t)) (\rho + \delta)\lambda_s(t) - \lambda'_s(t) = u(a(t)) + \beta a(t)A(t)N_i(t) (\lambda_i(t) - \lambda_s(t)) (\rho + \delta)\lambda_i(t) - \lambda'_i(t) = u(a(t)) - \gamma (\kappa(N_i(t)) + \lambda_i(t)) + \alpha \beta a(t)A(t)N_s(t) (\lambda_i(t) - \lambda_s(t)).$$

As usual for equilibrium, we then impose the conditions a(t) = A(t), $n_s(t) = N_s(t)$, and $n_i(t) = N_i(t)$. If $\alpha = 0$, this returns the equilibrium equations. If $\alpha = 1$ and κ is constant, this returns the equations describing the dynamics of the planner's solution. Intermediate values of α capture a degree of imperfect altruism.

In this formulation, we note one natural limit of altruism: We think it is unlikely that individuals view their own behavior as having an impact on the total number of infected people $N_i(t)$ and hence on the risk of death captured by κ . This is really an aggregate outcome, in contrast to the possibility that one person's social activity makes someone else sick, which an individual is more likely to believe that she can control.

7 Quantitative Exercises

7.1 Calibration

We calibrate the model at a daily frequency to US data on the 2020 Covid-19 outbreak. We offer some robustness with regard to the most important choices made in this section in Section 8.

To begin with, we set $\rho = 0.05/365$ to capture a 5% annual discount rate. In addition, we set $\delta = 0.67/365$, which implies an expected duration at which a cure is found of 1.5 years. We highlight that this jointly implies a model of heavy discounting relative to a standard economic model.

Next, we set $\gamma = 1/7$ such that the expected length of sickness lasts 1 week. We recognize that the average disease lasts longer, but it appears that few people are infectious and asymptomatic for longer than a week. Lauer et al. (2020) report a median incubation period for COVID-19 of five days and that 98% of people who develop symptoms after an exposure do so within 11.5 days. Below, we offer some robustness with regard to this choice.

We calibrate β for the model to capture data on the doubling time at the onset of the Covid-19 outbreak. Specifically, we target an initial daily growth rate of the stock of infected equal to 30%, consistent with a doubling time of approximately 3 days.⁵ From (2), we have that $\frac{N'_i(0)}{N_i(0)} = \beta - \gamma$, giving $\beta = 0.3 + \gamma = 0.443$. This implies a basic reproduction number of $R_0 = \frac{\beta}{\gamma} = 3.1$. Since there appears to be considerable uncertainty and disagreement about the value of R_0 we note that this strategy of backing out its value only relies on the expected duration of infectivity and aggregate data on doubling time in the early stages, two numbers that appear to be relatively well understood. However, some recent estimates suggest a higher value for R_0 and several authors work with a lower γ . We therefore offer a robustness exercise below where we use a higher value of R_0 and a correspondingly lower value of γ , while still hitting the initial 30% daily growth rate.

We work with the following period utility function,

$$u(a) = \log a - a + 1.$$
(15)

We think of the first part as the gross returns from social activity, in particular consumption,

⁵JHU's Center for Systems Science and Engineering reports a doubling time of 2–5 days in the US in the very early stages of the epidemic (Dong, Du and Gardner, 2020).

and of the second part as the cost associated with it. This gives an interior solution at $a^* = 1$ in a disease-free environment, with u(1) = 0.

We now turn to the cost of disease. We assume that the infection mortality rate for the disease is $\pi = 0.002$, independent of $N_i(t)$. This is lower than many estimates of the case mortality rate, but this smaller number is consistent with evidence that there are many undetected cases in real-world populations. Importantly, a constant death rate shuts down the health care externality which arises as health care deteriorates when many people are infected. As we show below, we still find that it is optimal to delay infections and typically optimal to avoid a high peak infection rate.

We think of the cost κ as equal to πv , where v is the value of a statistical life (VSL). We follow Greenstone and Nigam (2020) in assuming that v =\$10 Million for the US.⁶ Roughly speaking, this number is based on evidence that a typical individual would pay \$10,000 to avoid a 0.1% probability of death. With a discount rate of ρ , this is equivalent to paying a constant stream of $\rho \times$ \$10,000 to avoid this death risk, or equivalently \$1.37 per day. Compare this to US consumption per capita of approximately \$45,000 per year or \$123 per day, and we reach the conclusion that people would permanently give up over 1.1 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 (15) 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 x = 0.011, then v solves

$$\frac{\log(1)}{\rho} - 0.001v = \frac{\log(1 - 0.011)}{\rho}$$

This implies $v = \frac{-1000 \log 0.989}{\rho} \approx 80,000$ in our model units. Multiplying these together gives $\kappa = \pi v = 160$.

We note that this value is in the same ballpark as the one chosen in several other recent paper on the outbreak. For instance, Alvarez, Argente and Lippi (2020) choose a five times higher fatality rate π but a far lower VSL of \$1.3 million. Similarly, Hall, Jones and Klenow (2020) work with a VSL 50% higher than Alvarez, Argente and Lippi (2020) (and so still far below our value) but a four-fold higher death rate which implies a very similar value of κ . Finally, Eichenbaum et al. (2020) pick a VSL of \$9.3 million and a fatality rate of 0.5 percent, which implies a two fold higher κ . We offer a robustness exercise below where we consider doubling κ to reflect a higher infected fatality rate or a higher VSL.

Finally, we assume that initially $N_i(0) = 10^{-6}$. Prior to date 0, we assume $A(t) = a^* = 1$,

⁶See Greenstone and Nigam (2020) for a useful discussion of this value and the use of VSL in calculations such as ours.

so the disease grew without any response in social activity. From a very low initial prevalence, this implies that approximately $N_i(0)/(R_0 - 1)$ individuals have recovered or died by date 0, leaving the remaining 0.9999985 individuals susceptible.

Parameters			
Parameter Description	Parameter	Value	Target
Conditional transmission prob. Rate at which illness ends Cost of infection Arrival rate of cure Discounting Fraction initially affected	$egin{array}{c} eta \ \gamma \ \kappa \ \delta \ ho \ N_i(0) \end{array}$	$\begin{array}{c} 0.3+\gamma\\ ^{1/7}\\ 160\\ ^{0.67/365}\\ ^{0.05/365}\\ 10^{-6}\end{array}$	Initial doubling time Duration until symptomatic Death rate and VSL Exp. time until vaccine/cure Annual discount rate
Other			
Basic reproduction number Fraction initially susceptible	$\frac{R_0}{N_s(0)}$	$3.1 \\ 0.9999985$	Implied by γ and β no social distancing before $t = 0$

We summarize our calibration strategy in Table 1.

Notes: We calibrate the model at a daily frequency.

 Table 1: Calibration

7.2 Results

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.

7.2.1 Basic SIRD Model

Figure 3 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 bottom panel shows the share infected on a log scale. The pandemic unfolds rapidly even though only one out of one million individuals is initially infected. After several weeks, a sizable share of the population is infected. The infection rate peaks after seven weeks above 31 percent. As a consequence of the height of the peak, the benefits of herd immunity do not kick in before almost everyone is sick. By 14 weeks into the infection, only 5.3 percent of the population remains susceptible and 0.19% of the population has died, more than 600,000 people in a country the size of the United States. Although the pandemic ends quickly, the cost of the disease is substantial.



Notes: We set A(t) to its optimal level absent disease, $a^* = 1$ and use equations (1) and (2).

Figure 3: Basic SIRD Model.

Measured in utility units, it is -136.6, equivalent to a permanent reduction in social activity to a = 0.819.

Of course, this model completely fails to capture the experience in places that did not institute any restrictions on social activity. For instance, Sweden hit 1040 total cases on March 15, 2020.⁷ One month later, this number rose By mid April, this number had risen to about 11,927 confirmed cases despite the laissez-faire approach taken by the Swedish government. In contrast, our calibrated SIRD model predicts an increase by a factor of $e^{0.3\times30}$, or more than 8000-fold, during this period. Likewise, the SafeGraph micro-data document a remarkably uniform decline in individual social activity. The fact that this decline happened across the board in the US despite the large differences in policies also suggests that the basic SIRD model fails to capture a key aspect of this epidemic, namely that individual behavior responds to the risk of infection.

7.2.2 Laissez-Faire Equilibrium

We thus next turn to the disease dynamics in our laissez-faire equilibrium, which are depicted in Figure 4.⁸ The top two panels again depict the share of susceptible and infected. The

⁷Retrieved from https://www.worldometers.info/coronavirus/country/sweden/.

⁸Note that all the figures are conditional on no cure having been found.

difference between laissez-faire and the basic SIRD model is stark. Despite the government not intervening at all, the peak infection rate is one tenth of the level in the SIRD model, 3.5 percent. In turn, the response of individual behavior substantially prolongs the epidemic, with the infection rate staying elevated for a much longer time, albeit at a lower level. This implies that the population reaches herd immunity at a far lower level of N_i , compared with the SIRD model. Taking into account the possibility of a cure, in expectation just under 50 percent of people eventually get sick. Thus the expected death rate is about half as high as in the model without a behavioral response.

The third panel shows that individuals reduce their social activity by as much as 40 percent. The fourth panel depicts the ratio of the effective reproduction number to the fraction of susceptible people, $R_e(t)/N_s(t)$ on a double log scale.⁹ Recall that this is the number of newly infected individuals for each infected person which would prevail if everyone were susceptible for a given level of social activity A(t). This falls substantially, driving down the doubling time for the disease, but remains strictly above 1. As a consequence, the fact that N_i eventually starts falling is a consequence of the stock of susceptible people becoming smaller. Putting this together, the total welfare loss in equilibrium is equivalent to a permanent reduction in social activity to A = 0.854.

We note that the dynamics of social activity, A, under laissez-faire mirrors the behavior of N_i . The reason is simply that there is little private incentives to lower social activity when the risk of individual infection is negligible. As a consequence, individuals do not restrict activity until infections are rampant.

Taken together, the internalized part of the disease risk substantially delays the "wave," lowers the peak infection rate, and slows infections to an extent that allows the population to achieve a higher asymptotic susceptible share. In fact, as we show next, the laissez-faire equilibrium dynamics are closer to the optimal dynamics than to the SIRD model. This suggests that explicitly modeling the internalized dimension of disease outbreak is of first order importance.

7.2.3 Optimal Policy

Figure 5 contrasts the laissez-faire with the optimal policy. Reflecting the external effects of social activity, the key property of the optimal policy is delay. While peak infection in the SIRD model occur after 49 days and the equilibrium behavioral response delays the peak until 62 days have lapsed, the optimal policy delays it for 341 days. Because infections increase more slowly, the peak level of infection is also far lower under the optimal policy,

⁹On a double log scale, the vertical distance between two points y_1 and y_2 is proportional $\log(\log y_1) - \log(\log y_2)$.



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

Figure 4: Laissez-Faire equilibrium.

at 0.28 percent. We stress that this strong desire to "flatten the curve" is true in a model without any explicit cost of peak-loading of infection rates, i.e. where κ is constant. A health care externality would make the case for flattening the curve even stronger.

An optimal policy buys time for a cure. Asymptotically, a bit more than $1 - \gamma/\beta = 0.677$ of the population will eventually get sick, assuming a cure is never developed. This reflects the fact that activity will optimally asymptote back to a^* . Taking into account the possibility of a cure, however, reduces that to just 8.4 percent. Thus the expected death rate is one-sixth as large under the optimal policy as in equilibrium. The resulting welfare loss is equivalent to a permanent reduction in social activity to A = 0.907, significantly less than the equilibrium reduction to 0.854. The welfare cost is also well below the cost of permanently suppressing the disease by setting $R_e(t) = N_s(t)$, i.e. by setting $A = \sqrt{\gamma/\beta} = 0.568$ until a cure is found. The welfare cost of this policy is equivalent to permanently reducing social activity to A = 0.870.

The optimal policy achieves delay by acting preemptively. The third panel in Figure 5 shows the degree of social distancing in laissez-faire and optimum. At the outbreak of the disease, equilibrium behavior does not change because the risk of individual infection is negligible. Optimal policy, however, immediately curtails social activity. The planner recognizes that lowering the initial transmission rate buys time. In particular, the bottom panel shows that the ratio of the effective reproduction number to the share of susceptible individuals $R_e(t)/N_s(t)$ is far below its uncurtailed counterpart R_0 (which corresponds to the intercept of the laissez-faire time path). 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 $u'(a^*) = 0$.¹⁰

Despite the initial reduction in social activity, a remarkable feature of the optimal policy is that social distancing is never extremely intense. The planner could, of course, push the effective reproduction number $R_e(t)$ below the share of susceptible people $N_s(t)$, ending the disease, but he never chooses to do so, instead choosing values slightly above $N_s(t)$. While the stock of infected people $N_i(t)$ eventually starts declining under the optimal policy, this is a consequence of the fact that the stock of susceptible people $N_s(t)$ falls substantially below 1, combined with a limited reduction in social activity A(t).

To gain some intuition for this observation, we note that social activity eventually needs to return to its pre-pandemic level. The reason is that under any feasible policy, the share of infected individuals must converge to zero.¹¹ If A(t) were converging to a number smaller

¹⁰In contrast, Alvarez, Argente and Lippi (2020) assume the economy without disease is in a corner and so reductions in social activity, lockdowns in their words, have a first order cost.

¹¹Formally, we have that the sum of the number of recovered and deceased people evolve as $(N'_r(t) + N'_d(t)) = \gamma N_i(t)$ (equations 3 and 3). Since $N_r(t) + N_d(t)$ is bounded above by 1, their sum must converge,



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

Figure 5: Optimal Policy vs Laissez-Faire.

than $a^* = 1$, there would be a first order gain from a temporary and small increase in A, while the cost would be negligible if $N_i(t)$ is sufficiently small. Thus long-runs social distancing cannot be optimal.

Now suppose we suppress $R_e(t)$ below $N_s(t)$ at some early time t. Doing this for a while will reduce the infection rate to a negligible share of the population. But since the number of infected people never reaches zero, any attempt to relax social distancing will quickly lead to a reemergence of the disease, quickly undoing the effect of keeping $R_e(t)$ below $N_s(t)$. Thus setting $R_e(t)$ below $N_s(t)$ only makes sense if the intent is to keep this policy in place forever. But we have just argued that this is not optimal.

Finally, we discuss the shape of the recovery. We note that, 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 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. These restrictions do disappear eventually, with $R_e(t)/N_s(t)$ converging to R_0 as the level of the susceptible population falls slightly below $1/R_0$.

Overall, an important observation is that the planner achieves a delay in infections over the first year of the pandemic without completely locking down the economy. The key instead is an early and long-lasting reduction in social activity that is moderate in magnitude.

7.3 Revisiting the Evidence on Change in Individual Behavior

We briefly re-visit the quantitative evidence from SafeGraph from section 3. We note that none of the targets we chose in calibrating our model were actually related to the response of social activity to the Covid-19 outbreak. We do not take a stance on whether the hastily implemented lockdowns and mobility restrictions are 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 and 2 show a decline of some 25-50% in terms of activity across our three different metrics. We note that these metrics all have an inherent cardinality and so their decline is quantitatively meaningful. This is well captured by our laissez-faire model which suggests a decline of individual activity by 40% as can be seen in Figure 5.

Furthermore, we note that the model also captures the pace of the decline surprisingly well. In the model, equilibrium social activity declines form 98% on day 20 to 63% on day

which requires $N_i(t)$ converges to zero.

51. The contraction in the SafeGuard data was even faster. For example, Figure 1 shows that POI visits fell from 98 percent on March 11 to 45 percent on March 20 as state and local governments started to issue stay-at-home and shelter-in-place orders. The difference might reflect a delayed understanding of the seriousness of the disease.

We complement this analysis with an exploration of aggregate data available from Google at https://www.google.com/covid19/mobility/ for Sweden. This data uses Google's location history to track mobility along various dimensions relative to a baseline. Sweden has arguably been the Western country with the least restrictions on mobility and social activity throughout March 2020. The Swedish data currently cover the time period from February 23 to April 5.

We restrict attention to the categories "Retail & recreation," "Transit Stations," and "Workplace". The data suggest a reduction of activity of about 25%, 39%, and 28% for these three categories, respectively. This suggests that Swedes indeed internalize parts of the risk associated with the disease and therefore reduced activity, in line with our model. Like the US micro-data presented in section 3, the data squares up well with the laissez-faire dynamics under our preferred model calibration, both in terms of the size and pace of the decline.

We also note that we could have directly targeted this type of data in our calibration, in particular to select a value for κ . That is, a natural approach would be to let the individual response in behavior "reveal" the perceived cost of infection instead of relying on controversial and noisy direct measures of the infection mortality rate, π , and the value of a statistical life, v. However, as we show below, the responsiveness of the time path of A(t) to variation in $\kappa \equiv \pi v$ is limited and so we have opted with the direct measures. Nonetheless, we believe that the response of behavior absent policy intervention can be a useful source of information to reveal perceptions given the current level of uncertainty surrounding many key parameters.

8 Robustness

In this section, we show the laissez-faire and optimal dynamics given perturbations to various parameters. Given the large uncertainty around many of the key modeling parameters we consider fairly large changes. We find that our main findings are robust to alternative parameter choices. In particular, we find that 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 with one notable exception, a limit in the extent of social distancing such that the effective reproduction number $R_e(t)$ stays above the share of susceptible individuals $N_s(t)$.

8.1 Alternative Utility Function

In Figure 6, we modify the utility function to $u(a) = -\frac{1}{2}(1-a)^2$. This leaves u(1) = u'(1) = 0and also leaves u''(1) 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 above the share of susceptible individuals.

8.2 Cost of Disease

This subsection shows the results for $\kappa = 320$, that is we double our baseline parameter value of the expected cost of infection. For instance, this can be viewed as capturing a death rate of $\pi = 0.004$ (compared with $\pi = 0.002$ in the baseline). We show the resulting dynamics for optimum and laissez-faire in Figure 7.

Doubling the cost of infection cuts the peak equilibrium infection rate by more than half, from 3.3 percent to 1.5 percent. The peak optimal infection rate declines by a larger percent, although from a very low level of 0.276 percent to a negligible 0.014 percent. The optimal policy delays a sizable outbreak for an extremely long time, effectively until a cure has been found with very high probability. It does so by initially reducing social activity and then continuing to suppress it to a level with $R_e(t)$ just above $N_s(t)$ for decades. While not visible in the figure, social activity is eventually allowed to return to normal and a fraction less than γ/β of individuals remain disease free. We note of course that all of this again is conditional on no cure being found, an exceedingly unlikely scenario.

8.3 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 8. 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



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

Figure 6: Optimal Policy vs Laissez-Faire with quadratic utility.



Notes: See Table 1 for calibration. Only change: We set $\kappa = 320$ instead of $\kappa = 160$. Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 7: Optimal Policy vs Laissez-Faire with high cost of infection.

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 $R_e(t) > N_s(t)$, as before.

8.4 Lower Cure Probability

We have so far assumed that the expected time until a cure is 1.5 years. Figure 9 shows what happens if we double this to 3 years. The laissez-faire equilibrium dynamics are effectively the same, reflecting the fact that 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, only briefly interrupted by a more severe constraint on activity still with $R_e(t) > N_s(t)$ —at the peak of the infection. In particular, social distancing lasts even once it appears that the share of susceptible people has reached a plateau.

8.5 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 the 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 10.



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). Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 8: Optimal Policy vs Laissez-Faire with long duration of infectivity.



Notes: See Table 1 for calibration. Only change: We set $\delta = \frac{0.33}{365}$ (instead of $\frac{0.67}{365}$). Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 9: Optimal Policy vs Laissez-Faire with low probability of cure.

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 for the first 67 days to such an extent that $R_e(t)$ falls below $N_s(t)$; this cannot be shown on the double log scale. 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 turn, remains largely unchanged. Starting with more infections just shifts it ahead in time.

9 Next Steps

9.1 Heterogeneity

Several key dimensions of heterogeneity come to mind. In particular, a key feature of the Covid-19 pandemic is the case fatality rate by age: Individuals age 65 and older seem to have a far higher fatality rate than younger ones. At the same time, one might argue that they have a smaller cost associated with a reduced social activity since they are primarily retired.

We believe that our current framework can naturally and easily be extended to accommodate these dimensions of heterogeneity. With k types of people, we would need to solve a system of 2k differential equations in equilibrium. Assuming a utilitarian social planner that can separately choose the level of social activity for each type, there would be a corresponding system of 2k differential equations describing the optimal policy. We are currently working on that extension.

9.2 Antibody Tests

We have assumed that everyone knows when their infection ends. In reality, few people are currently tested, and so many recovered people cannot be sure whether they were sick. We can capture this by separating the recovered state into two categories, depending on whether the individual is aware they were sick. Someone who does not know they have recovered from Covid-19 must choose the same level of social activity as the susceptible and infected, given their lack of information. In this environment, antibody testing may be useful for letting recovered people know that they are recovered. This has the dual benefit of allowing those individuals to return to a normal level of social activity and of discouraging social activity by those who are still susceptible or infected.



Notes: See Table 1 for calibration. Only change: We set $N_i(0) = 10^{-2}$ instead of $N_i(0) = 10^{-6}$. This also reduces the initial value of $N_s(0)$ to 0.985 instead of 0.9999985. Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 10: Optimal Policy vs Laissez-Faire with higher stock of initially infected.

9.3 Vaccines

We have so far assumed that at a random future date, a cure will end all costs associated with the disease. This is an unlikely scenario. More plausible is the gradual roll-out of a vaccine that shifts some people from susceptible to recovered without enduring an infection. To the extent the vaccine provides imperfect coverage, even a 100 percent vaccination rate will leave some people vulnerable to future infections, and potentially slow the recovery in social activity.

10 Conclusion

This paper uses standard dynamic optimal choice tools from economics to integrate privately optimal behavior and policy analysis into an epidemiological model. Our quantitative exercises reveal several robust patterns: Even in laissez-faire, individuals sharply reduce social activity due to risk of infection gradually. However, an optimal policy immediately curtails social activity immediately and delays the full outbreak to buy time. But even the optimal policy lets the disease eventually run its course: If no cure is found a large number of people eventually get infected. The framework we develop is general and tractable: Optimal behavior and policy are encoded in an additional set of differential equations that can jointly be solved with the epidemiological block. We therefore view the tools offered here as a natural building block to explore the role of various additional features of Covid-19 and their interplay with optimal policy.

References

- 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.
- **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.
- 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.

- 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 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.
- **Piguillem, Facundo and Liyan Shi**, "The optimal covid-19 quarantine and testing policies," Technical Report, Einaudi Institute for Economics and Finance (EIEF) 2020.